NASA Final Report

System Safety for Highly Distributed Air Traffic Management

Award Number NNXI4AC71A

Nancy Leveson, Cody Fleming, Cameron Thornberry MIT

Chris Wilkinson, Honeywell Aerospace Advanced Technology

March 31, 2015

Abstract

This research has two main objectives. The first objective is to develop rigorous, systematic tools for the analysis of future concepts in order to identify hazardous scenarios and undocumented assumptions. Related to this theme is the problem of assessing safety-related risk when little design detail is available, with the goal of assisting concept development and design when modifications are most effective. The second objective is to extend these tools to assist stakeholders in the development of concepts using a safety-driven approach. Ideally, this safety-guided concept development would supplement existing system engineering activities, including architectural and design studies that occur during tradespace exploration. Both objectives especially apply to systems where the tradespace includes human operation, automation or decision support tools, and the coordination of decision making agents

Keywords

Early concept safety analysis, safety assessment, safety assessment, risk assessment, hazard analysis, formal model, trajectory based operations, NextGen, safety management, human factors, safety directed design

Executive Summary

The next generation of air traffic management systems will involve significant changes from the way ATC (air traffic control) is done today. Reliance on software is increasing and allowing greater system complexity. Humans are assuming supervisory roles over automation, requiring more cognitively complex human decision-making. Control is shifting from the ground to the aircraft and shared responsibilities. In addition, coupling and interconnection between land, airborne, and space systems introduces more potential for accidents stemming from unsafe and unintended component interactions.

Traditional hazard analysis and risk management techniques, most of which were created 50 or more years ago for the much simpler systems of that time, cannot effectively handle the more complex systems being developed today. More powerful hazard analysis methods are needed.

Systems-Theoretic Early Concept Analysis (STECA) is a new type of analysis technique based on a very different type of paradigm and assumptions about the causes of accidents. Traditional approaches to safety assume accidents are caused by component faults and failures. Those approaches therefore focus on component reliability enhancement and do not adequately handle accidents caused by unsafe and unintended interactions among non-failed components. Such unsafe interactions usually stem from system design errors, software requirements errors, or operator errors.

In addition, traditional approaches to safety treat human error as random or stochastic. Human error, and all human behavior, is affected by the context and system design in which it occurs. For purposes of improving safety and designing to reduce errors, it is more useful to provide information about *why* the error occurred, which is needed to eliminate or mitigate errors in the system design, than to assign a probability of error.

The STECA approach is based on a new accident causality model called STAMP, for Systems-Theoretic Accident Model and Processes. STAMP treats safety as a control problem rather than a component reliability problem. Because of its basis in systems theory, STECA can handle the advanced features of NextGen and the complexity of the proposed operational improvements. It provides a rigorous process to assist subject matter experts, in other words, it allows an "organized inquiry." There is potential for automated assistance in the analysis.

Most current hazard analysis techniques are usable only late in the development process, when the major design decisions have already been made and a concrete design is available to be analyzed. However, in addition to analyzing and verifying an existing design, designers of complex systems need assistance and analysis techniques during early concept development. Seventy to ninety percent of safety-related design decisions are made before most hazard analysis techniques are applicable [Frola and Miller, 1984]. STECA works early in concept design so that decisions can be evaluated before they are very costly to change. It identifies potential causal scenarios leading to system hazards and the system and subsystem design requirements to mitigate or control these potential design flaws.

Systems-Theoretic Early Concept Analysis (STECA) is demonstrated in this report as a potential alternative or addition to the traditional hazard analysis techniques. The demonstration is done on an important component of NextGen called Trajectory-Based Operations (TBO). TBO is gate-to-gate spacing and trajectory optimization enhancement that specifies aircraft trajectories in four dimensions (three dimensions in space, and time-of-arrival at specified points in space).

Because TBO involves advanced decision making capability both from the ground and on the flight deck, the TBO analysis identifies unsafe control actions of flight crews in addition to en route air traffic controllers. The STECA analysis identifies causes of inconsistencies, lack of coordination, unenforced safety constraints, and conflicts among goals. Based on these findings, requirements are then written and allocated to system components to prevent these causes. Requirements are allocated to surveillance systems, traffic monitoring systems, ground- and flight deck-based automation, air traffic controllers, pilots, trajectory modeling systems, and weather monitoring and prediction systems. Requirements are also allocated to ensure appropriate coordination between these components.

STECA includes non-linear, indirect, and feedback relationships among events and actors, and therefore many of the scenarios identified in this report include more than just component failure. Therefore many of the requirements identified by the STECA analysis go beyond reliability and relate to the behavior of system components (both human controllers and automation) and the information that those components receive and exchange. The hazard analysis identified several scenarios and/or causal factors that are not considered in the TBO Concept of Operations [JPDO 2011] and related hazard analyses that use traditional approaches [JPDO 2012].

The causal factors identified by STECA but not included in the existing TBO documentation include potential lack of coordination between controllers both within and across sectors, lack of coordination and understanding between ground and airborne elements, timing of TBO clearances originating from different control agents, potential lack of synchronization and coordination between surveillance and monitoring sources, and conflicts between TBO automation and other tools and Air Traffic Management tasks.

These additional scenarios and causal factors can be used for future revisions of the TBO ConOps as well as in the development of alternative designs or architectures and guide future trade studies. This report thus includes a section devoted to identifying and developing alternative architectures that eliminate or mitigate certain hazardous scenarios.

While the results by themselves will help stakeholders make more informed decisions with respect to delegating authority, responsibility, procedural control, and algorithmic control within the NAS, extensions to the STECA analysis done in this report could assist in ascertaining relative levels of risk associated with different design and control decisions. For example, the paradigm used in the current NAS is hierarchical and centralized relative to the future vision of NextGen. Alternatively, aspects of TBO are highly de-centralized. A hierarchical, centralized structure with clearly delineated control authority might be more appropriate for tactical (short-term) resolutions. On the other hand, decentralized control may be more appropriate during strategic (long-term) negotiations. The different types of hazard causes that result from changes

in control structure, different time horizons (or timing requirements), and goals could be compared. STECA should then be used as a part of future safety-guided design efforts as further technical details become available.

One of the greatest strengths of the STAMP accident model is that it can capture detailed causal factors related to controllers, software, component interactions, and overall system *design*. In addition to these systemic causal factors, the treatment of the controller (whether human or automated) also analyzes the impact that feedback and internal models of the system states have on the controller, and how flaws in these could lead to an accident. While existing STAMP-based analytical tools have come far in improving the analysis of human control agents, this report presents an extension to improve the techniques of generating and categorizing causal factors related to the human controller and to integrate concepts from within the field of Human Factors.

Finally, this report assesses the role of STAMP, STECA, and related techniques in the overall safety management framework advocated by the FAA. STECA does not require changes in the current Air Traffic Organization (ATO) Safety Management System and fits within it in three places. It could be used to identify hazards in the Safety Risk Management (SRM) pillar. It also provides a useful way to describe systems in the first step of the SRM process. The functional control structure generated to perform STPA includes all the system aspects currently described for this SRM descriptive step as well as a more complete description of system functional behavior than that currently described as a minimum for SRM. Finally, it could be used in the Safety Assurance pillar to identify what needs to be assured, including audit procedures to verify that safety is not degrading as behavior changes over time. And, of course, because STECA works early in the system design process, it can be used to assist in developing and finding gaps in the ConOps and other design documentation.

Table of Contents

1. Research Goals, Objectives and Plan

Current preliminary hazard analysis and risk assessment techniques are limited with respect to the kinds of scenarios they identify and how risk is communicated to decision makers. Instead of using traditional failure-based approaches to analyze a concept, this report uses a different approach. The proposed approach to analyzing and developing a concept is based on control- and systems-theory that identifies more hazardous scenarios and assists stakeholders in the development and refinement of a concept. By identifying more hazardous scenarios with limited design information, decision makers can eliminate or mitigate hazards by the selection of appropriate architectural options when the cost of doing so is much less than when a design is nearly complete.

Therefore, this research has two main objectives. The first objective is to develop rigorous, systematic tools for the analysis of future concepts in order to identify hazardous scenarios and undocumented assumptions. Related to this theme is the problem of assessing safety-related risk when little design detail is available, with the goal of assisting concept development and design when modifications are most effective. The second objective is to extend these tools to assist stakeholders in the development of concepts using a safety-driven approach. Ideally, this safety-guided concept development would supplement existing system engineering activities, including architectural and design studies that occur during tradespace exploration. Both objectives especially apply to systems where the tradespace includes human operation, automation or decision support tools, and the coordination of decision making agents.

The next generation of air traffic management (called NextGen in the U.S. and SESAR in Europe) will include increased coupling and interconnectivity among airborne, ground, and satellite systems and intensive use of computers and software in safety-critical roles. Control will be shifting from the ground to the aircraft and to shared responsibility for safety among ATC (air traffic control), pilots, and airline operations centers.

The planned coupling and interconnection between land, airborne, and satellite systems introduces more potential for accidents stemming from unsafe and unintended component interactions.

NextGen's success will lie in the program's ability to design new ATM (Air Traffic Management) operational increments (OI) that do not decrease the safety of the current system. One of the requirements for achieving this goal is the ability to assess the safety of proposed changes. There are several current approaches being used or tried for NextGen. All involve the use of fault trees, event trees, bow-tie diagrams, Bayesian belief networks, Markov models, and other techniques to perform probabilistic risk analyses (PRA), simulations, mockups and demonstrations, etc. The foundational analysis techniques for the assessment (fault trees and event trees) are all 40 to 50 years old and predate the extensive use of computers in complex systems. They do not handle the level of complexity underlying the plans for ATM upgrades and cannot be extended to do so because they are based on an underlying accident causal model that excludes the planned extensions in coupling and interconnectivity. The traditional hazard analysis techniques also are limited to completed designs and thus are not very useful in designing safety into the system from the beginning.

Our hypothesis is that these techniques miss the most important types of problems that are going to occur in a computer-intensive ATM system and that a totally different type of safety assessment paradigm is required. In this research, we investigated an innovative and very different type of hazard analysis, based on systems theory rather than the traditional reliability theory underlying the other methods.

The goal of this research is to (1) extend and demonstrate our new risk assessment paradigm for the problems that need to be solved to implement NextGen and (2) compare and evaluate the results with the current (traditional) hazard analysis techniques currently being used for NextGen. To accomplish this goal, we developed a new hazard analysis method to guide safety-driven design, applied it to a future NextGen OI and compare the results to the risk assessment process currently being used by the FAA. The objectives are listed in Table 1 along with the report section where the objective is addressed.

Table 1. Objectives/Report Cross Reference

Obj	Description	Section
1	Create a hazard analysis technique, based on STAMP, that is useful in the early concept design phase of NextGen operational increments and can be used to guide decision making by engineers (subject matter experts) as they define the new procedures. The approach should not only be able to analyze the ConOps (Concept of Operations) produced in the early stages of design (which STPA can now do) to find safety flaws, but it will be able to assist the designers in producing a ConOps with fewer potential safety flaws. Time and money will be saved by better decision-making in early design rather than incurring the enormous costs of rework later in the development process.	5
2	Devise a method to compare the risk involved in various alternative NextGen architectures.	7
3	Demonstrate how STPA-SDD (now called STECA) can be used to derive verifiable safety requirements and to derive scenarios for use in simulation experiments.	7
5	Evaluate how the new STPA-SDD would fit into the current FAA Safety Management System	11
6	Extend the current analysis approach for human error in STPA to account for more types of human factors problems such as flawed decision-making, external factors such as distraction, etc. While much automation is to be added in NextGen, humans will still play critical roles in the anticipated changes for quite a while. STPA is better than the current hazard analysis methods, such as fault trees and event trees, in realistically treating human error. But much more is needed.	9

7	NextGen will involve phased rollout and mixed fleets, integration of UAVs, etc. Evaluate how much work will be involved in performing STPA-SDD on changes to systems that were previously designed using STPA-SDD and compare this cost with the amount of effort in doing a more traditional risk analysis on upgrades. Starting from scratch each time is not practical. How much safety analysis rework is necessary to integrate changes to ATM?	5, 6, 7
8	Compare the results of STPA-SDD with the current approach to safety analysis being used on highly-distributed, tightly coupled systems such as NextGen.	8

To meet those objectives, the following work plan was devised.

(1) Select a stretch goal of NextGen on which to develop and demonstrate our new analysis technique (by end of Month 1)
(2) Develop STPA-SDD and apply it to the demonstration system. This development included:
  (a) Create of a way to compare architectures with respect to safety,
  (b) Define of a system engineering process that integrates STPA-SDD into the standard system engineering process,
  (c) Generate detailed safety requirements for the demonstration system,
  (d) Extend of STPA to include sophisticated human factors analysis
  (e) Develop leading indicators for the demonstration system
(3) Evaluation of the new approach
  (a) Compare the STPA-SDD results with the standard risk assessment techniques being used for NextGen today on the demonstration system
  (b) Analyze how STPA-SDD would fit into the current FAA Safety Management System
  (c) Evaluate the amount of analysis rework necessary to make changes to the demonstration system in the future

2.   Motivation for This Work

Often the perception among engineers and other stakeholders is that safety is expensive. Safety-related features are also seen as intrusive because they seem to result in reduced performance, increased weight, or unnecessary complexity. In fact safety often is costly, both in terms of economics and technical performance, but this is not due to any intrinsic property of safety itself. Rather, the reason safety costs so much is that it is often considered only after the major architectural tradeoffs and design decisions have been made. Once the basic design is finalized, the only choice is to add expensive redundancy or excessive design margins [Leveson, 2009].

It has been estimated in the defense community that 70-80% of the decisions affecting safety are made in the early concept development stages of a project [Frola and Miller, 1984]. As Figure 1 illustrates, compensating later for making poor choices at the beginning can be very costly. As Stieglitz [1948] said, "Safety must be designed and built into airplanes, just as are performance, stability, and structural integrity". Stieglitz' quote is appropriate for all complex systems, not merely airplane design. Safety must be designed and built into systems from the very beginning of concept development.

Unfortunately, traditional tools used for analyzing and improving safety are only applicable in the later stages of system development, when detailed design information is available. These same tools were developed long ago, when the primary cause of accidents was due to mechanical failure [Vesely et al., 1981]. Modern systems exhibit hazardous behavior due to a series of factors that extend well beyond hardware failure. The introduction of new technology, such as computers and software, is changing the types of accidents we see today [Leveson, 1995; Leveson, 2012].

Hazardous behavior arises in systems due to unsafe interactions between components, even when the components have not necessarily failed. Given the complexity of today's systems, these interactions are increasingly difficult to understand and predict. The underlying assumptions of traditional hazard analysis tools also oversimplify the role of human operators [Dekker, 2005; Rasmussen, 1997; Woods et al., 2010] and software requirements errors [Leveson, 2009; Lutz and Carmen Mikulski, 2003]. Not only are traditional hazard analysis techniques incapable of analyzing systems that are immature in terms of design detail, they are also very limited with respect to these new accident causation factors, which will become increasingly prevalent in tomorrow's systems.
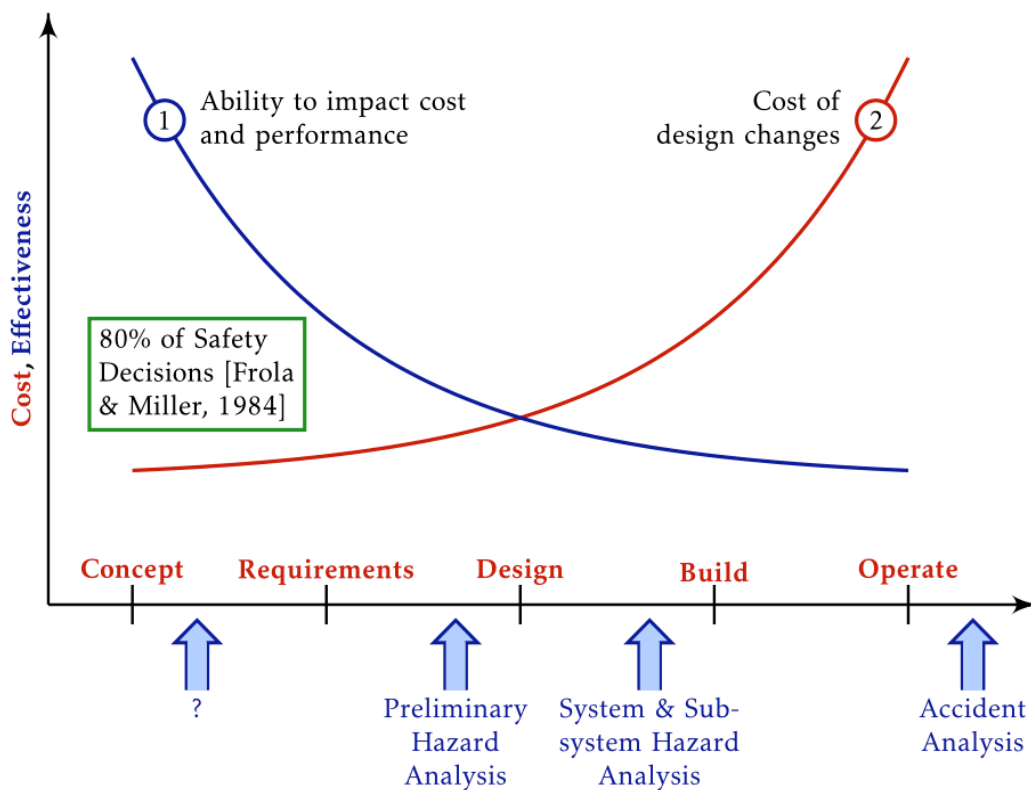
Figure 1. Decision Effectiveness during Life Cycle (adapted from (Strafaci, 2008))

The current[1] national airspace (NAS) system in the United States has achieved historically low accident rates, with an exponential decrease in major[2] accidents since 1960 (Boeing, 2009).

---

[1] "Current" refers here to the system as it exists (existed) before NextGen implementations.

However, the forecasted growth in passenger and freight flights is expected to be more than 5% in the coming decades [Netjasov and Janic, 2008], and the current air traffic management system cannot sustain this growth. In addition to increasing capacity demands, the United States national airspace faces increasing pressure to reduce greenhouse gas emissions and operator costs.

To meet these challenges, the Federal Aviation Administration (FAA) has developed a program called NextGen, which "integrates new and existing technologies, policies and procedures to reduce delays, save fuel and lower aircraft exhaust emissions to deliver a better travel experience" [FAA, 2013c]. The proposed changes must also maintain or improve the FAA's stated top priority, ensuring safe skies and airfields. The European counterpart to the FAA, EUROCONTROL, faces similar challenges and has an analogous program to NextGen called Single European Sky ATM Research (SESAR) [Patteau, 2009].

NextGen calls for increased focus on more efficient flight paths, a shift in responsibility from ground-based crews to flight crews and their flight deck-based decision support tools, the use of trajectory-based operations instead of clearance-based maneuvering of aircraft, and many other changes. In short, NextGen will result in increased reliance on automation, greater coupling between ground and aircraft technology, a major shift in the way airspace information is gathered and disseminated, and a total revamping of how aircraft paths are managed. All of these changes will come as the result of incremental upgrades that span years and even decades.

The FAA and other institutions involved in NextGen recognize that these changes pose a risk to safety-related properties of the current national airspace. Fortunately, these institutions have also recognized the importance of understanding safety-related risk early in development of NextGen improvements as well as including safety during concept development and design selection activities [FAA, 2012c; JPDO, 2012]. Unfortunately, however, many of the existing techniques are limited with respect to these goals [Harkleroad et al., 2013].

While the FAA and EUROCONTROL are faced with new issues in air traffic management, aircraft manufacturers have their own challenges. Aircraft manufacturers deal with customers who increasingly weigh cost factors as well as improvements in performance, comfort, and environmental effects [Croft, 2005]. As a result, Boeing developed the B787 and B737-MAX, Airbus developed the A380 and A350, and smaller companies like Embraer and Bombardier also continue to innovate. Boeing selected lithium cobalt oxide (LiCo) batteries for its B787 power system in order to save weight while meeting power requirements. Multiple incidents with these batteries caused the entire B787 fleet to be grounded [Brown, 2013]. The problem was extremely difficult to fix due to constraints on space and weight, taking months of design work and thousands of hours of testing to complete [Yeo, 2013].

Airbus has had its own problems, as Qantas grounded its entire A380 fleet after an uncontained engine failure [ATSB, 2013]. One suggestion was to reduce the allowed thrust on the engines, but this was deemed unprofitable [Yeo, 2010]. Ultimately the engines had to be replaced [Kollewe and Gabbatt, 2010].

The issues with the B787 and A380 highlight the difficulty of identifying and mitigating hazards in complex systems with highly coupled sub-systems and components. The issues also show how

---

[2] Major accidents, per the Boeing study, include those with fatalities and/or hull loss.

difficult, expensive, and sub-optimal it is to modify an existing design in order to mitigate safety-related issues.

Other aerospace domains face similar challenges. The United States ponders returning to the moon and sending astronauts to Mars while the European Union, Japan, China, India, and Russia attempt to advance their human spaceflight programs [Drake et al., 2010; Irvine, 2007]. As the United States National Aeronautical and Space Administration (NASA) learned throughout its space shuttle program, early architectural decisions have significant and sometimes irreversible implications for safety. As just one example, all previous US manned space vehicles had launch escape systems, and such a system was discussed for the shuttle program. However, NASA did not implement a launch escape system for programmatic and technical reasons [McCurdy, 1993]. After the Challenger disaster, the shuttle orbiters were fitted to allow for crew evacuation, though this could only be used when the shuttle was in a controlled glide. The crew would have had to reach the exit from their seats and jump out, and this solution is irrelevant for launch events [Petty, 2002; Dumoulin, 1988]. Like the Boeing and Airbus problems, the shuttle program illustrates the importance of safety considerations during early architectural decisions and the diminishing returns after design decisions have been made.

This report is not intended to be a critique of recent aircraft designs or of the shuttle program. However, the problems Airbus, Boeing, and NASA have experienced highlight the fact that early design decisions have significant implications for safety and can have major cost impacts later in the program (Figure 1). The importance of integrating safety analysis into early systems engineering activities cannot be overemphasized.

3.   Background to Current Aviation Safety Assessment Practices

Much of the literature regarding safety during early system engineering activities focuses on the development of preliminary hazard lists, preliminary hazard analyses, and the assessment of those results. Safety is rarely included explicitly or analytically in traditional early system engineering activities such as concept generation, concept selection, and tradespace exploration.

Preliminary hazard analysis (PHA) is one of the earliest safety-related activities, and many of these analyses not only attempt to assess risk but also suggest potential mitigations or design modifications. However, as Section 3.1 describes, PHA is still limited with respect to influencing design and ensuring that safety-related properties are integrated into the system as early as possible. Recall from Section 2 the importance of considering safety in the design process, when design modifications have minimal cost and maximum influence.

3.1   Safety Activities during Preliminary Design

Traditionally, safety-related activities conducted during the preliminary phases of an engineering program include developing Preliminary Hazard Lists (PHL), performing Preliminary Hazard Analysis (PHA), and informing decision-makers by using risk assessment techniques. The primary objective of a preliminary hazard list is to identify very high level hazards during concept development, while the objective of preliminary hazard analysis is to categorize the risk level (hazard level) of the identified hazards.

Comparisons of system architectures and design alternatives are based on trade studies that incorporate performance objectives such as (for aerospace systems) mass, speed, range, and

efficiency, as well as cost estimates. While PHA efforts typically begin before an architecture is selected, these safety efforts are performed in parallel with architecture studies and therefore have little impact the general systems engineering process. That is, the PHA is totally independent of the architectural study. In addition, the focus on component failure inherent in current PHA techniques severely restricts the ability to identify hazard causality due to human behavior, software requirements flaws, and the interaction between human operators and software.

PHA is a guided effort for identifying hazards, their associated causal factors, effects, and level of risk. Mitigating measures are also sometimes included [Ericson, 2005]. Currently and in the past, PHA has focused on failure modes of sub-systems or components, but such a focus is limited for several reasons.

First, there are many causes of hazardous behavior other than component failures or faults, which is explored further in later sections. Second, PHA starts during concept formation when little design detail is available. This lack of design detail, combined with the novelty of many systems being developed today, makes it infeasible to assign a likelihood and/or consequence to a component failure. Before describing the preceding limitations in greater detail, current PHA guidance and their underlying assumptions must be discussed.

3.1.1 Preliminary Hazard Analysis (PHA)

Standard preliminary hazard analyses include a list of hazards to be avoided, potential causes of those hazards, effects on the system, severity level of the hazards, and supporting comments or recommendations [Vincoli, 2005]. Table 2 shows a generic PHA table and expected contents. Most of the guidance for conducting a PHA comes from military or regulatory bodies, although there are notable exceptions in the academic and practitioner literature. This sub-section explores both government and academic sources.

Military and aerospace standards provide further guidance on expected inputs and outputs for a preliminary hazard analysis. The United States military standard for system safety includes a relatively systematic procedure for identifying hazards, classifying mitigation measures, and documenting all of the above. The U.S. military specifies that the PHA shall identify hazards by considering the potential contribution to mishaps from: system components, energy sources, hazardous materials, interfaces and controls, software, human factors engineering and human error analysis, and several other factors [DoD, 2012]. Other military standards provide similar guidance for conducting a preliminary hazard analysis, see for example [MoD, 1996].

Table 2. Sample PHA Worksheet, adapted from [Vincoli, 2005]

PRELIMINARY HAZARD ANALYSIS						
PROGRAM: _____				DATE: _____		
ENGINEER: _____				PAGE: _____		
ITEM	HAZARDOUS CONDITION	CAUSE	EFFECTS	RAC	ASSESS-MENTS	RECOMMEN-DATIONS
Assigned number sequence	List the nature of the condition	Describe what is causing the stated condition to exist	If allowed to go uncorrected, what will be the effect or effects of the hazardous condition	Hazard Level assignment	Probability, possibility of occurrence	Recommended actions to eliminate or control the hazard

The United States Military Standard 882E [DoD, 2012] additionally requires that a risk assessment be performed for all the factors identified in the PHA. General aspects of risk assessment during preliminary phases of system development are described elsewhere [Fleming 2015].

The United States Federal Aviation Administration prescribes a five-step process for performing a risk assessment and recommends several analytical processes or techniques for identifying hazards and causes [FAA, 2008]. Operational Hazard Assessment (OHA) and Comparative Safety Assessment (CSA) are touted as the most relevant tools for hazard analysis during concept development or preliminary design. Both OHA and CSA share characteristics with efforts labeled PHA throughout this sub-section.

OHA contains a list of operational hazards[3], their effects, and severity classifications for each item. Like the other PHA outputs described in this sub-section, CSA includes a severity and likelihood classification for each hazard, cause, and effect. CSA is different from many of the other references in that it is intended for use when design or operational changes are proposed for the national airspace. Therefore, the results of CSA are meant to be compared with some baseline system.

Roland and Moriarty [2009] suggest a slightly different set of activities that should be included in PHA. While the artifacts of the analysis remain the same (hazards, causes, likelihood, consequence), Roland and Moriarty suggest specific activities and aspects of the system that should be scrutinized. The PHA should consist of a review of historical safety experience in similar systems; an examination of basic energy sources; an examination of safety-related interfaces; exposure to environmental hazards such as shock, vibration, extreme temperatures; an examination of software modules in their interfaces with hardware, operators, or other software for possible hazards; and safety related equipment including interlocks, redundancy, and fail-safe designs.

Ericson [2005] describes the necessary inputs, desired outputs, and expected process of a PHA. To perform a PHA, the analyst must have (1) a preliminary hazard list, (2) design knowledge, and (3) hazard knowledge. Preliminary hazard lists are developed even earlier in a project's development than a PHA and are typically the first safety-related assessment of a project. To possess design knowledge, Ericson emphasizes the importance of having a list of system components and their intended function. Hazard knowledge is derived from hazard checklists and consists of a basic knowledge of hazards, their sources and components (element, initiating mechanism, and target), and hazards in similar or heritage systems. According to Ericson, hazard checklists should include:

> (1) Energy sources
> (2) Hazardous functions
> (3) Hazardous operations
> (4) Hazardous components
> (5) Hazardous materials
> (6) Lessons learned from similar type systems

---

[3] Operational Hazards in the FAA's OHA framework are equivalent to hazard *causes* in the other PHA references.

(7) Undesired mishaps

(8) Failure mode and failure state considerations

Figure 2 depicts the mapping of top-level mishaps (TLM) to system safety requirements (SSR) and safety critical functions (SCF) via a decomposition of top-level mishaps into lower-level failure modes. In addition to outlining what should be expected in terms of inputs, outputs, and process, Ericson provides several real-world PHA examples as well as a list of common mistakes to avoid.

Although Ericson provides more comprehensive guidance and documentation than the other literature reviewed here, the underlying theory and expected outputs of all of the preceding literature are mutually consistent. PHA is intended to generate a list of hazard causes arising from faulted or failed components, assign a probability of occurrence along with severity of consequence, and then (combined with potential mitigation measures) assess the relative risk of each cause.
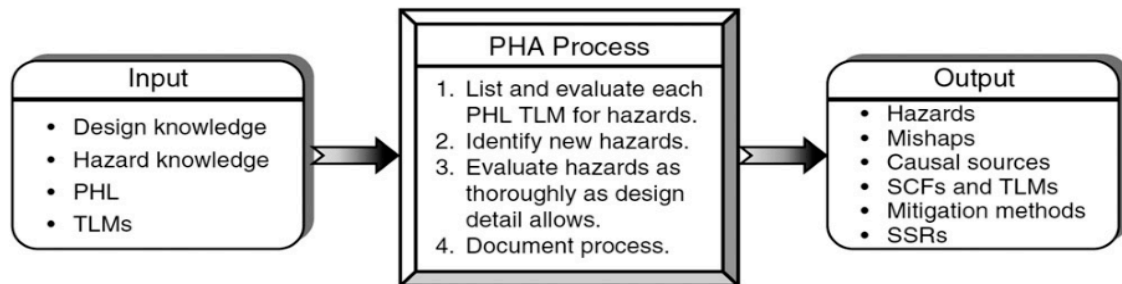
Figure 2. PHA Inputs, Process, and Outputs [Ericson, 2005]

For illustrative purposes consider a recent example of PHA from the aerospace domain. Table 3 includes a subset of the PHA results for trajectory-based operations (TBO), a proposed upgrade in the NextGen program. Observe the emphasis on component failure, likelihood of occurrence, and significance of failure in Table 3. The emphasis on component failures in this analysis is linked to an accident causality model, although this link is not explicitly stated in any of the literature on PHA. The next section describes and compares accident causality models.

In summary, PHA has historically been a framework for identifying hazard causes in terms of component and sub-system failures or faults. In addition to their effects on system hazards, these causes are further identified in terms of their severity of consequence and probability of occurrence.

Table 3. PHA for Trajectory-Based Operations, adapted from (JPDO, 2012)

Hazard ID	Hazard Name	Hazard Description	Causes	Signifi-cance	Likeli-hood	Assumed Mitigations	Strength of Mitigations	Outcome Risk	Justification
TBO-0004	ADS-B Ground System Comm Failure	GBA does not receive ADS-B message	Receiver failure	High	Low	Redundant equipment; SSR; Primary Radar; Overlapping ADS-B coverage; Multi-Lat; Design and Equipment Certification Requirements	Medium	Medium	Strength of Mitigations depends on the type of backup; Multi-lat should be used if spacing requirements are tighter than they are today. ...
TBO-0021	GBA fails to recognize dynamic situation and is unable to find a solution	The software lacks robustness in its implementation that leads to inability to find a solution	Design flaw, coding error, insufficient software testing, software OS problem	High	Med	Comprehensive system testing before certification and operational approval. TCAS; See and avoid. Pilot could recognize in some cases; Controller could recognize in some cases	Low / Medium	Medium / High	Anything that is complex can lead to this situation

3.1.2 Limitations in Current Approaches to Preliminary Hazard Analysis

Many accidents in modern, software-intensive systems involve the interaction of components, the incorrect specification of software requirements, and unsafe human behavior due to a myriad of factors including human errors caused by confusion resulting from the automation design [Leveson, 2012]. Safety is an emergent property; that is, safety (or lack thereof) only emerges from the interactions of a system's components. The next section describes emergence in greater detail.

Much of the guidance for identifying hazard causes during PHA—for example documenting energy sources, hazardous materials, or faulted (failed) modes of mechanical components—is a necessary but insufficient aspect of hazard analysis. The underlying model of accident causation in the PHA literature is reductionist and assumes that if all component failure modes of a system have been identified, then so have all potential sources of hazardous behavior. Current PHA literature provides little to no guidance for how to identify hazardous interactions amongst components; incorrectly specified software requirements; or human operator errors due to poor design of procedures, computer interfaces, and underlying logic of automation and decision support tools. One could list "component interaction error", "software design error", or "operator mode confusion" in a hazard list like Table 2. However, there is no guidance on how to identify why these factors might occur. That is, simply listing "software error" without further explanation does not add any useful information to the engineering process.

Consider again the example from Ericson [2005]. Table 4 lists the potential failure modes for the missile rocket booster subsystem of a fictitious missile system. PHA-4 lists "erroneous initiate commands", with software faults and human error as potential causes. The recommended action is then to use multiple switches or conduct a fault tree analysis of the fuze design. Though the example PHA correctly identifies software and human operators as potential sources of hazardous behavior, the method is limited in helping the analyst to reason about specific causal factors[4]. That is, all hazards are caused by hardware, software, or humans. Simply listing a generic set of factors is not very helpful, and PHA techniques suffer from a lack of guidance in identifying causal factors that lead to specific hazardous states that stakeholders wish to avoid.

Table 4. ACE Missile Example (adapted from [Ericson, 2005])

Subsystem: Missile Warhead Subsystem				
No.	Hazard	Causes	Effects	Recommended Action

---

[4] This is not a critique the specific example in [Ericson, 2005] but is a critique of PHA methods in general.

PHA-4	Inadvertent W/H explosives initiation due to erroneous initiate commands	Erroneous commands from hardware faults; software faults; human error	Personnel death/injury	Use multiple independent switches in fuze design Conduct FTA of fuze design
PHA-5	Inadvertent W/H explosives initiation due to external environment	Bulletstrike, shrapnel, heat	Personnel death/injury	Use insensitive munitions
PHA-6	Failure of W/H explosives to initiate when commanded	Hardware faults; software faults	Dud missile; **not a safety concern**	

Finally, PHA requires the analysts to assign a probability of occurrence, which could be either a probability of a faulted mode, the probability of its affect on the system, or both. There are several problems with this approach. First, PHA occurs near the beginning of a project when very little design detail is available. Often the analyses rely on heritage data to extrapolate a probability for identical or similar components. However, in the case of software and human operators, heritage data does not, and cannot, exist. Even when software modules are reused, either the context in which they will operate or some details of the code will be modified. Of equal importance, human and software behavior is not stochastic.

Software behaves exactly as it is specified and coded [Leveson, 1995], and human behavior is highly dependent on context, both in the system design and due to environmental factors [Dekker, 2005; Leveson, 2012; Klien et al., 2004; Reason, 2000; Vicente, 1999]. Even in later stages of system development, or during operations, software error probabilities cannot be calculated and human error probabilities can only be attained with limited applicability.

In summary, the types of causes identified by PHA techniques are limited to electro-mechanical faults or very generic causes related to human or software behavior. These generic types of causes are not particularly useful for guiding the design. In addition, when PHA requires probability estimates, the probabilities of many sources of hazardous behavior are difficult or impossible to validate. The probabilities are either highly uncertain (due to limited design detail at the beginning of a program), unknowable (due to human behavior in a complex environment), or irrelevant (because software behavior is strictly deterministic).

4.   Systems Theoretic Accident Model and Process

System-Theoretic Accident Model and Processes (STAMP) was created to capture more types of accident causal factors including social and organizational structures, new kinds of human error, design and requirements flaws, and dysfunctional interactions among non-failed components [Leveson, 2004; Leveson, 2012]. Rather than treating safety as a failure problem or simplifying accidents to a linear chain of events, STAMP treats safety as a control problem in which

accidents arise from complex dynamic processes that may operate concurrently and interact to create unsafe situations.

Accidents can then be prevented by identifying and enforcing constraints on component interactions. This model captures accidents due to component failure, but also explains increasingly common component interaction accidents that occur in complex systems without any component failures. For example, software can create unsafe situations by behaving exactly as instructed or operators and automated controllers can individually perform as intended but together they may create unexpected or dangerous conditions.

STAMP is based on systems theory and control theory. In systems theory, emergent properties are those system properties that arise from the interactions among components. Safety is a type of emergent property. The emergent properties associated with a set of components are related to constraints upon the degrees of freedom of those components' behavior [Checkland, 1999]. There are always constraints or controls that exist on the interactions among components in any complex system. These behavioral controls may include physical laws, designed fail-safe mechanisms to handle component failures, policies, and procedures. Such controls must be designed such that the safety constraints are enforced on the potential interactions between the system components. In air traffic control, for example, the system is designed to prevent loss of separation among aircraft.

System safety can then be reformulated as a system control problem rather than a component reliability problem—accidents occur when component failures, external disturbances, and/or potentially unsafe interactions among system components are not handled adequately or controlled, leading to the violation of required safety constraints on component behavior (such as maintaining minimum separation). System controls may be managerial, organizational, physical, operational, or in manufacturing. In STAMP, the safety controls in a system are embodied in the hierarchical safety control structure. The next section describes hierarchy theory in greater detail, while Figure 3 shows a generic, example hierarchical control structure.
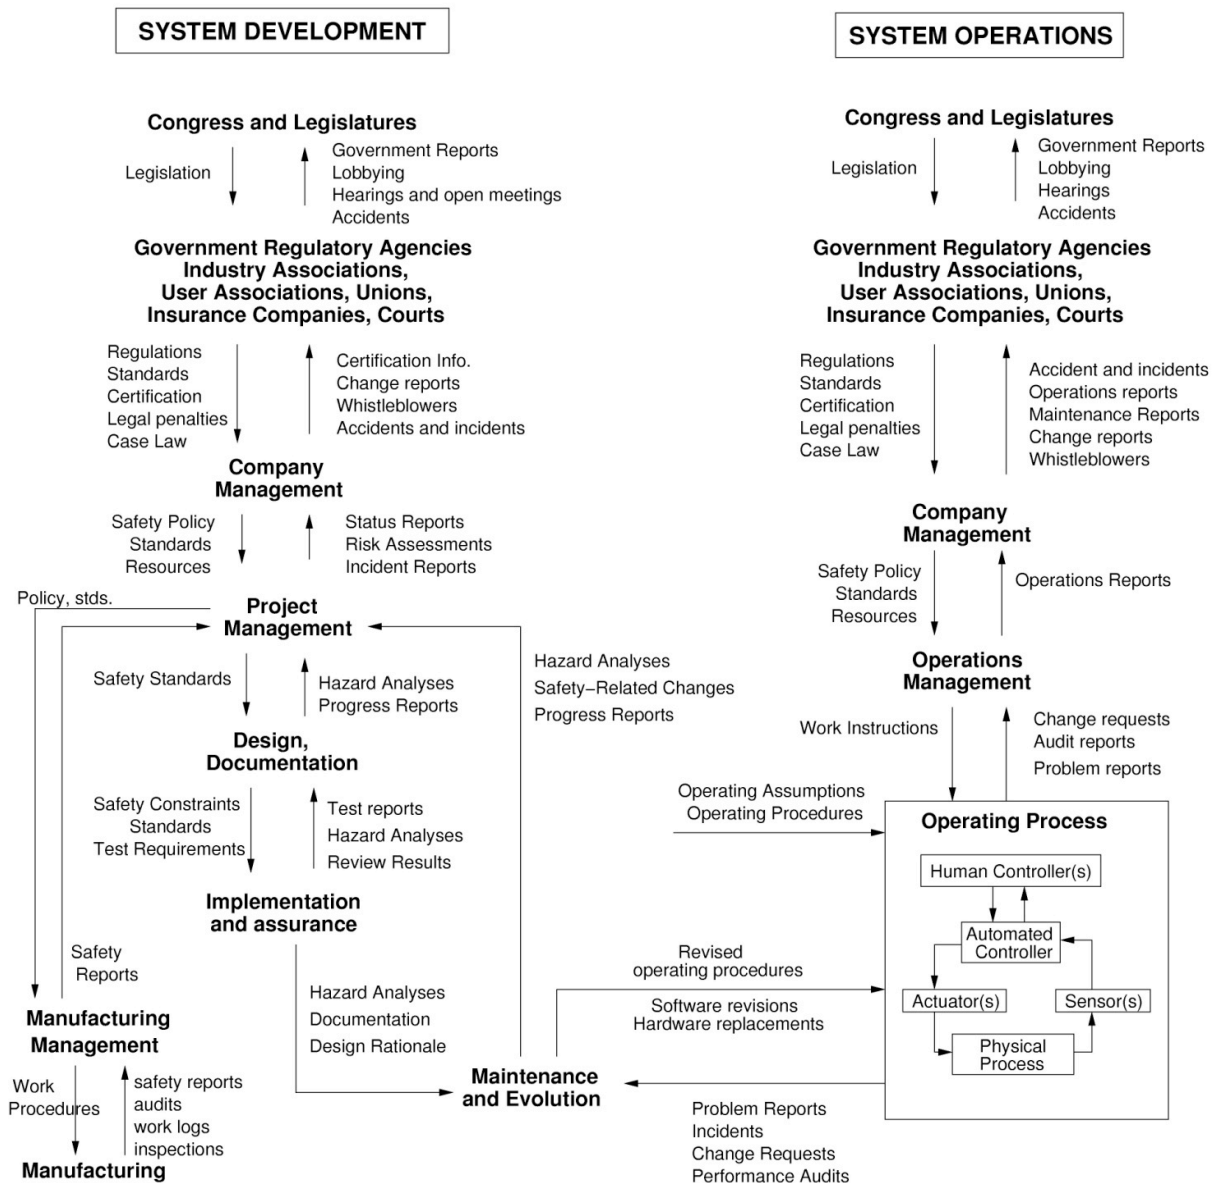
Figure 3. An Example Safety Control Structure

Control processes operate throughout the hierarchy whereby commands or control actions are issued from higher levels to lower levels and feedback is provided from lower levels to higher levels (see Figure 3 and Figure 6). Accidents arise from inadequate enforcement of safety constraints, for example due to missing or incorrect feedback, inadequate control actions, component failure, uncontrolled disturbances, or other flaws. STAMP defines four types of unsafe control actions that must be eliminated or controlled to prevent accidents:

- Not providing the control action causes the hazard
- Providing the control action causes the hazard
- The timing or sequencing of control actions leads to the hazard

- The duration of a continuous (non-discrete) control action, i.e., too short or too long, leads to the hazard.

One potential cause of a hazardous control action in STAMP is an inadequate process model used by human or automated controllers. A process model contains the controller's understanding of 1) the current state of the controlled process, 2) the desired state of the controlled process, and 3) the ways the process can change state. This model is used by the controller to determine what control actions are needed. In software, the process model is usually implemented in variables and embedded in the program algorithms. For humans, the process model is often called the "mental model" [Leveson, 2004]. Software and human errors frequently result from incorrect process models. Accidents like this can occur when an incorrect or incomplete process model causes a controller to provide control actions that are hazardous. While process model flaws are not the only cause of accidents in STAMP, it is a major contributor.

The generic control loop in Figure 4 shows other factors that may cause unsafe control actions. Consider an unsafe control action for an air traffic controller: a flight crew is instructed to increase altitude while another aircraft is flying through that new altitude. The control loop in Figure 4 would show that one potential cause of that action is an incorrect belief that the airspace above the aircraft is clear (an incorrect process model). The incorrect process model, in turn, may be the result of inadequate feedback provided by a failed sensor, the feedback may be delayed, the data may have been corrupted, etc. Alternatively, the system may have operated exactly as designed but the designers may have omitted a feedback signal or the feedback requirements may be insufficient.

Control input or external information wrong or missing

Controller

Inadequate Control Algorithm (Flaws in creation, Process changes, Incorrect modification or adaptation)

Process Model inconsistent, incomplete, or incorrect

Inappropriate, ineffective or missing control action

Process Model inconsistent, incomplete, or incorrect

Inadequate or missing feedback Feedback delays

Actuator
Inadequate Operation

Sensor
Inadequate Operation

Delayed operation

Controller 2

Conflicting control actions

Controlled Process

Component failures
Changes over time

Process input missing or wrong

Process output contributes to hazard

Unidentified or out-of-range disturbance

Incorrect or no information provided
Measurement inaccuracies
Feedback delays

Figure 4. STPA Control Loop with Causal Factors

4.1.1 Techniques Based on STAMP

Systems-Theoretic Process Analysis (STPA) is a hazard analysis technique that is based on the STAMP accident model. Causal Analysis based on STAMP (CAST) is used for accident analysis, and STPA-Sec is used for security analysis. STPA is significantly more powerful than failure-based techniques in the ability to capture a wider array of hazardous behaviors, including organizational aspects, requirements flaws, design errors, complex human behavior, and component failures [Leveson, 2012]. While many hazard analysis techniques stop once a sequence of events or failures has been identified, STPA helps explain the complex reasons why a sequence of events might occur, including underlying processes and control flaws that may exist without any component failure.

Although STPA is relatively new compared to traditional methods, it has been demonstrated successfully on a wide range of systems including aviation [Fleming et al., 2013], spacecraft [Ishimatsu et al., 2010; Nakao et al., 2012; Fleming et al., 2012], missile defense systems [Pereira et al., 2006], civil infrastructure [Dong, 2012], and others.

STPA has only been applied to existing, operational systems or to projects with a significant amount of design detail, although Harkleroad et al. [2013] identified it as a potentially effective method during concept development. STPA has been most effectively applied when the actions available to a control agent are discrete or when an agent's available actions are pre-specified. In

addition, STPA has not been used to compare architectures or design tradeoffs in terms of safety-related risk.

Figure 5 shows the analytical tools that are based on the STAMP accident causality model. In addition to STPA, CAST is an accident investigation tool based on STAMP, and STPA-Sec is a new technique used to identify and control vulnerabilities in the security domain (Young and Leveson, 2014). These tools are especially adept at capturing behavior in modern complex human- and software-intensive systems where component interaction accidents (or security incidents) have become increasingly common and traditional chain of events models are inadequate. STECA, the new concept development technique described in the next section, is also based on the STAMP accident causality model. These tools then support more general processes, such as systems engineering, management, operations, and regulation.
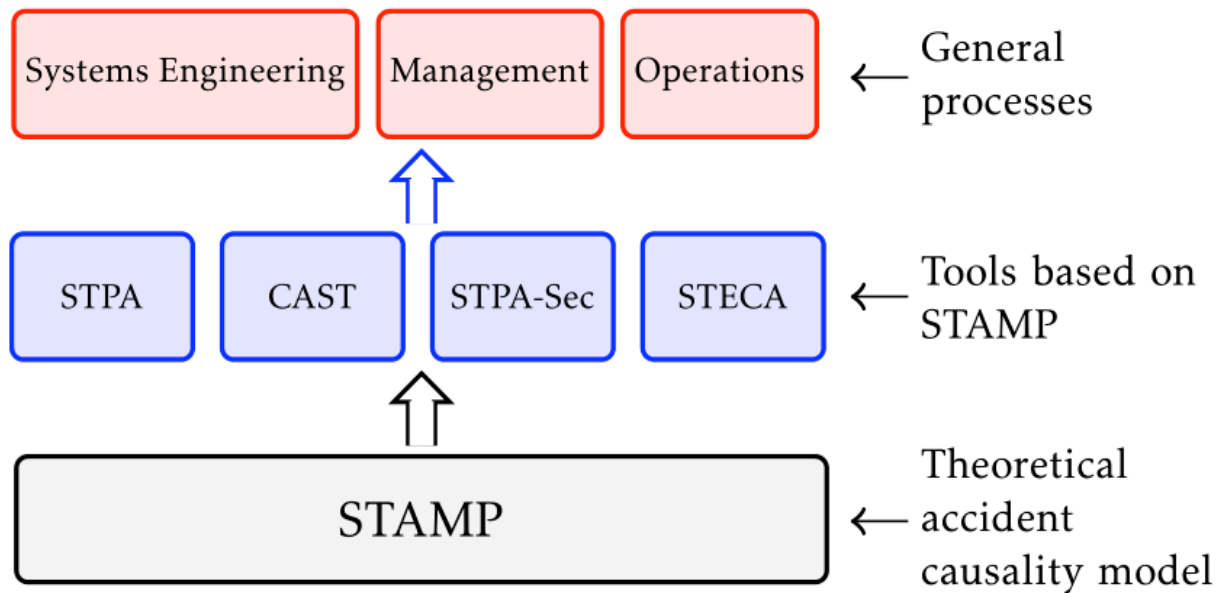
Figure 5. Techniques based on STAMP Accident Causality Model

4.1.2   General Assessment of Accident Models and Hazard Analysis Techniques

Tomorrow's aerospace systems will be highly coupled, involve nonlinear dynamics, and require changing roles among human operators and software systems. The hazard analysis must be able to handle these issues. Additionally, the motivation described in Section 2 calls for a method that can be applied early during concept development and system architecting, when little design detail is available and analytical techniques can help guide system development.

There are two general (and related) limitations to many of the traditional hazard analysis techniques, which include Fault Tree Analysis (FTA), Failure Modes and Effects Analysis (FMEA), and Event Tree Analysis (ETA). Although attempts have been made to extend some of these methods to include interaction and dependence among causal factors, the methods are still limited by their underlying accident causality models. This chain-of-events causality model assumes that accidents are caused by a linear succession of discrete, failure-based events and necessarily omits feedback and interaction elements of accident causality.

Moreover, these methods came along when engineered systems were much different than today. The use of software and computer systems has exploded since the 1960s, and although many attempts have been made to adapt these methods the original assumptions persist.

FMEA is a good method for analyzing component reliability, while FTA is a more efficient hazard analysis technique because it considers only those failures that lead to the higher level events. ETA is only effective when events proceed in a consistent, chronological order. None of the above techniques are particularly effective during concept development and tradespace exploration, when little design detail is available.

STPA has many advantages over the traditional techniques due to its underlying model of accident causation and the guidance it provides analysts in identifying hazardous scenarios. Because of its demonstrated ability to identify software design and requirements flaws; potentially hazardous, context-dependent human behavior; and hazardous interaction among components, STPA has the most potential for successful application to complex aerospace systems. In addition, STPA relies on a functional model of the system, as opposed to a physical model, which provides further rationale for its application in the early phases of development. STPA will need to be extended in order to handle the problems inherent in concept development and system architecting, which represents the motivation for developing the new approach described in Section 5.

4.1.3   Summary

The initial safety-related activity in most aerospace projects, the preliminary hazard analysis, is limited by an accident causality model that pre-dates the development of computer-intensive systems. The guidance given for conducting preliminary hazard analysis is restricted to consideration of mechanical failures or oversimplifies the role of human operators and software in modern, complex systems.

While progress has been made in generating safety-related requirements during certain systems engineering activities, little has been done to include safety in other activities such as concept generation and tradespace exploration. In fact, given the current body of tradespace exploration literature it may be inappropriate to include safety analytically (currently), due to the necessity of directly measurable metrics in many of the tradespace exploration frameworks.

A new approach is needed to overcome these difficulties. This report introduces a new process, based on the STAMP model, for analysis of a concept of operations to assist in safety-driven design from the early stages of system engineering.

5.   Systems-Theoretic Early Concept Analysis (STECA)

Section 1 outlined two broad objectives for this research. The first objective is to develop rigorous, systematic tools for the analysis of future concepts in order to identify hazardous scenarios and undocumented assumptions. The second objective is to extend these tools to assist stakeholders in the development of concepts using a safety-driven approach. Both goals especially apply to systems where the tradespace includes human operation, automation or decision support tools, and the coordination of decision making agents.

In order to improve upon the existing state of practice and theory of including safety during concept development, a methodology should help analysts and stakeholders to *systematically*:

- identify *missing information* or *undocumented assumptions* that will be required for effective operation of the system;
- identify *inconsistent* or *conflicting* information within a concept that may lead to hazardous behavior;
- identify where *more specific* operational concepts are required to understand safety- and functionally-related behavior of the system;
- identify *requirements* or safety *constraints* for 1–3; and
- identify the *mitigation strategies* associated with factors identified in 1–3

Because of this report' emphasis on introducing systematic methods into earlier phases of systems engineering, it is important to note a general distinction between the above objectives. The first three objectives (1–3 above) involve developing and applying more sophisticated techniques in order to improve upon the current limitations of PHA, which lack systematic guidance and focus on component failure. The latter two objectives (4–5) extend to slightly different phases and goals of system engineering. Generating safety-related requirements and identifying mitigation strategies should explicitly be part of the early systems engineering effort and result in a process called *safety-guided design* and development.

To put these objectives in perspective, consider the timeline in Figure 1 from the beginning of this report and the following question: what is needed to move from a high-level Concept of Operations to later stages of systems engineering? By identifying hazardous scenarios in the ConOps and using this information to generate requirements, the approach can assist stakeholders and engineers in developing a system architecture.

Before describing the model-based approach to achieving these objectives, it is important to develop the theoretical underpinnings of the approach.

5.1   Theoretical Foundations

This report extends the System Theoretic Accident Model and Process (STAMP) and its associated hazard analysis technique (STPA – Systems-Theoretic Process Analysis) for analysis of an existing ConOps. The reasons for selecting this accident model and its associated analytical processes include: (1) during early system engineering activities, STAMP's focus on the functional behavior of a system makes it a strong candidate relative to other techniques that rely on analyzing physical hardware [Harkleroad et al., 2013]; (2) STAMP's ability to identify component interactions, software design flaws, and potential sources of hazardous human behavior, which are prevalent in many future systems under development in the aerospace industry; and (3) the systems- and control-theoretic underpinnings of the STAMP accident causality model can potentially be extended into more rigorous, formalized techniques that can guide early system engineering efforts.

The focus of this report requires further explanation of some of the general characteristics that are often present during concept development. The primary artifact of concept development, the ConOps often consists of natural language text and/or low fidelity graphical depictions of work- and information-flows. ConOps should contain some reference to stakeholder goals and system- level requirements, but they rarely contain specific design requirements.

Because a ConOps is developed long before the system becomes operational, it typically includes (often implicit) assumptions about the future that are not necessarily true when the

document is being developed. For example, a ConOps may contain assumptions about future technologies that do not yet exist. Finally, often in practice a ConOps is developed by committee, with disparate members who have different goals and perspectives [JPDO, 2011]. All of these characteristics make it difficult to systematically analyze a Concept of Operations and in particular to rigorously, systematically achieve Objectives 1-3 in Table 1.

5.1.1   Systems Theory, STAMP, and STPA

STAMP and STPA are explained briefly in Section 3, and the proposed extension rests on two related principles of the STAMP model of accident causality; (1) emergence and hierarchy and (2) communication and control. In fact, these two principles form the basis of general systems theory.

In systems theory, a level of complexity is characterized by properties that do not exist at lower levels [Checkland, 1999]. These properties are called emergent, and the study of any property that cannot be accounted for at lower levels of complexity is referred to as the theory of emergence. Safety is an emergent property, not a property of individual components.

For example, consider a pilot and an air traffic controller. One of the primary objectives of a pilot is to ensure that the aircraft follows a stable trajectory. Alternatively, from the perspective of an air traffic controller, who is tasked with managing (part of) the national airspace, the aircraft following a stable trajectory does not by itself constitute safety. From the level of the national airspace system, safety only emerges when the air traffic controller coordinates multiple aircraft trajectories and considers terrain, airspace restrictions, weather, aircraft capability, and other factors.

Hierarchy theory is concerned with the fundamental differences between one level of complexity and another. In the airspace example, at one level of complexity an air traffic controller must manage the trajectories of all the aircraft in his or her sector. At another level, a pilot simply cares about following his or her assigned trajectory. The pilot then must manage lower level characteristics such as thrust and control surface trim, which when combined will help in achieving that trajectory. Hierarchy theory provides an account of the relationships between different levels and how hierarchies are formed [Checkland, 1999]. That is, what separates the levels and what links them?

In a hierarchy of open[5] systems, such as the national airspace, maintaining the hierarchy involves processes where there is *communication* of information in order to *control* lower levels. Checkland [1999] states that a system requires communication and control if it is "to survive the knocks administered by the systems' environment". Returning to the air traffic control example, if the air traffic controller is tasked with ensuring safe separation between aircraft, (s)he must have means to control or impose constraints on individual aircraft behavior. The air traffic controller must also be provided information about aircraft states, for example if aircraft have to change altitude due to turbulence. Therefore, *imposing constraints* (control) plays a fundamental role in the proposed approach in this report, as does *feedback* (communication).

Four conditions are required for process control [Ashby, 1957; Leveson, 2012]:

---

[5] An open system is one that exchanges material, energy, and information with its environment.

(1) Goal condition: the controller must have a goal or goals
(2) Action condition: the controller must be able to affect the state of the system, typically by means of an actuator or actuators.
(3) Model condition: the controller must contain a model of the system
(4) Observability condition: the controller must be able to ascertain the state of the system, typically by feedback from a sensor

Figure 6 depicts a hierarchical system, where the components interact with their environment, one level imposes constraints on the level below it, and feedback about its performance is transmitted back to the level above it. The proposed approach in this report recognizes that safety is an emergent property. System components interact with the environment and with each other, and safety is enforced by a set of control laws or goal conditions that constrain the behavior of these components. With safety viewed as a control problem, accidents occur when component failures, external disturbances, and dysfunctional interactions among components are not adequately controlled [Leveson, 2012]. This view of accident causality forms the basis of STAMP, and STPA is the hazard analysis approach based on the STAMP model.
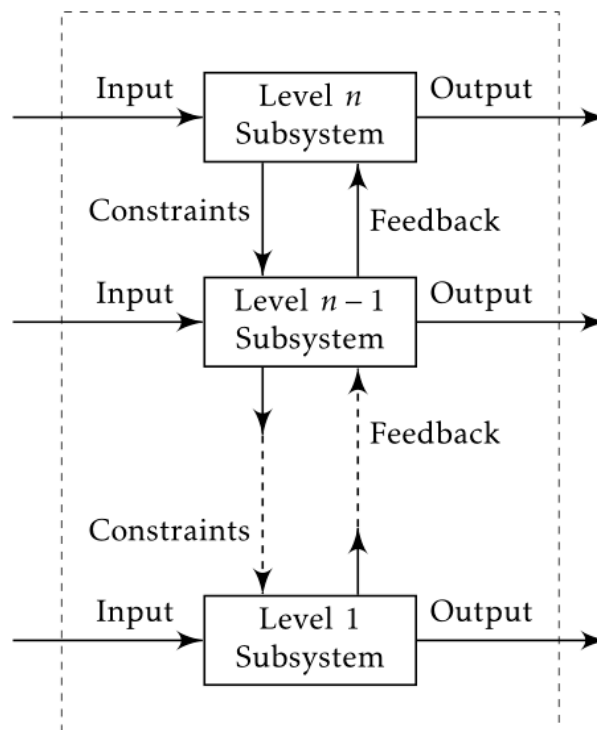
Figure 6. Basic Features of a Hierarchical System (adapted from [Mesarovic et al., 1970])

5.1.2   Outline of Approach

These concepts—hierarchy and emergence, and communication and control—are fundamental to the model-based approach proposed in this report. For the purposes of analysis, these concepts should be used in the opposite order. Control-theoretic concepts are used first to construct a model of the system, and theories of hierarchy and emergence (in addition to control and communication) are then used to interrogate the model itself. The process is conducted according

30

to Figure 7, where the main contributions from this extension are represented by the lower four boxes. The following sub-sections describe the theoretical development as well as provide a brief example for illustrative purposes.

Like a typical STPA hazard analysis, the systems-theoretic early concept analysis (STECA) begins with accidents and hazards, a high level decomposition of control functions, and then a set of high level safety responsibilities. These are basic system and safety engineering activities that should be done for any project (first box, Figure 7). Section 6 provides an example of how to identify a hierarchical list of safety responsibilities that is based on systems theory.

GENERAL,
SYSTEMS-THEORETIC
CONOPS ANALYSIS

SAFETY-DRIVEN DESIGN

ConOps

Identify System Hazards
• Demonstration: Sec 6.2

Derive System Safety Constraints

• Demonstration: Sec 6.2

Identify Control Concepts
• Description: Sec 5.2
• Demonstration: Sec 6.3

Derive Refined Safety Constraints

• Description: Sec 5.4
• Demonstration: Sec 7

Identify Hazardous Scenarios and Causal Factors

• Completeness Criteria: Sec 5.3.1
• Analyze Safety Resp: Sec 5.3.2
• Coordination & Consistency: Sec 5.3.3
• Demonstration: Sec 6.4

Refine, Modify Control Structure

• Description: Sec 5.4
• Demonstration: Sec 7

Figure 7. Proposed Methodology - STECA

31

5.2   *Systematic Control Model Development*

Potential benefits of model-based systems engineering include the use of repeatable processes, promoting consistent views of the system, and formal application of modeling to support requirements generation, design, analysis, and verification [Friedenthal et al., 2007]. It is in this vein that this research seeks to develop ConOps in terms of models rather than informal documentation.

Recall that the objectives of this research are to identify missing information or undocumented assumptions that will be required for effective operation of the system; identify inconsistent or conflicting information within a concept that may lead to hazardous behavior; and identify where more specific operational concepts are required to understand safety- and functionally-related behavior of the system. Consider also that a ConOps typically contains natural language text and graphical depictions of operating concepts, neither of which contain specifications nor rigorous, formal accounting of roles and responsibilities. The following modifications of STPA are needed to allow an analyst to rigorously and systematically develop a system model based on the descriptions contained in a ConOps.

Consider the STPA causal factors described in Section 3 and included again here in Figure 8. These guide-words are very effective for analyzing complete or nearly complete designs or specifications, and for identifying potential flaws in a system. However, a control-theoretic approach can also be used for understanding and specifying how a system should behave. Recognizing the roles that components of the control loop play in enforcing safe behavior can help in developing and analyzing a concept, where detailed design information is unavailable and using the guide-words in Figure 8 may not yet be the most effective.

Control input or external information wrong or missing

Controller

Inadequate Control Algorithm (Flaws in creation, Process changes, Incorrect modification or adaptation)

Process Model inconsistent, incomplete, or incorrect

Inappropriate, ineffective or missing control action

Inadequate or missing feedback
Feedback delays

Actuator
Inadequate Operation

Sensor
Inadequate Operation

Delayed operation

Incorrect or no information provided
Measurement inaccuracies
Feedback delays

Controller 2

Conflicting control actions

Controlled Process

Component failures
Changes over time

Process input missing or wrong

Process output contributes to hazard

Unidentified or out-of-range disturbance

Figure 8. STPA Control Loop with Causal Factors

5.2.1 Operational Roles in a Control-Theoretic Framework

Rather than directly using the control flaws (causal factors) from Figure 8, first examine the basic functions of each entity in the control loop. That is, what is required of each entity in the control loop for effective, safe system behavior? What are the responsibilities of the controller, actuator, controlled process, and sensor? How do these entities interact with each other, with the environment, and with other control loops?

The *Controller*:

- creates, generates, or modifies control actions based on algorithm or procedure and perceived model of system
- processes inputs from sensors to form and update process model
- processes inputs from external sources to form and update process model
- transmits instructions or status to other controllers or entities in the system

The *Actuator*:

- translates controller-generated action into process-specific instruction, force, heat, torque, or other mechanism

The *Controlled Process*:

- interacts with environment via forces, heat transfer, chemical reactions, or other input
- translates higher level control actions into control actions directed at lower level processes (if it is not at the bottom of a control hierarchy)

The *Sensor*:

- transmits continuous dynamic state measurements to controller (i.e. measures the behavior of controlled process via continuous or semi-continuous, digital data)
- transmits binary or discretized state data to controller (i.e. measures behavior of process relative to thresholds; For example, sensor has algorithm built-in to determine a threshold but has no control authority)
- synthesizes and integrates measurement data (e.g. takes location data from different types of sensors to create an estimate, like a Kalman filter)

This information can be built into a template that analysts and stakeholders use when developing, analyzing, and discussing a concept of operations. In fact, this information can be formalized into a formal, mathematical model that can be rigorously queried to ensure completeness and consistency. Such a formalization will be shown in the next sections.

The roles of the controller, actuator, controlled process, and sensor, and their interactions with the environment and other control loops can be summarized with 15 generic keywords or guide words. Figure 9 depicts these guide words in the familiar control loop format. With a proper accounting of these 15 items, the control loop can achieve the four necessary conditions[6] of process control and adequately interact with its environment, other processes, and other controllers. In other words, these guide words are necessary to ensure that a control loop is controllable and coordinable with other controlled processes.

---

[6] See section 3.1.1.
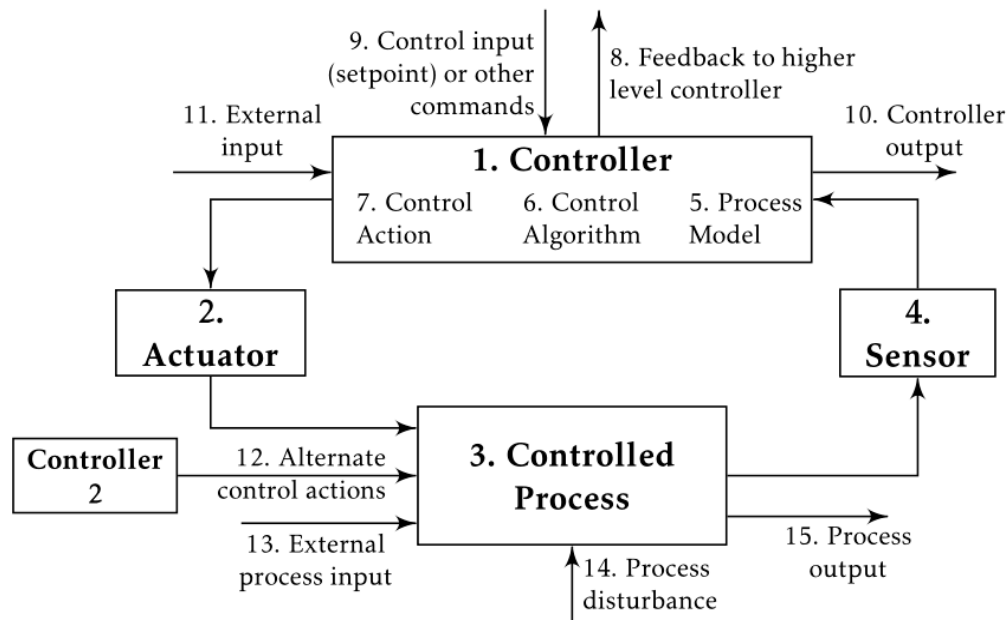
Figure 9. Control Loop with generic entities

The information in Figure 9 and the above lists (Controller, Actuator, Controlled Process, Sensor) can then be used to systematically parse and query the natural language description or graphical depiction in a concept of operations. The resulting model and subsequent database are easy to interrogate and visualize. These qualities help the analyst to check for internal inconsistencies and/or missing information that may result in unsatisfied control conditions, and also to check for inconsistencies across the system hierarchy.

Table 5 provides a series of prompts that an analyst can use when reading a text or graphic in a ConOps.

Table 5. Control-theoretic Analysis of Text

Source / Subject	What is the primary subject of the text? What is the primary source of action that the text (or graphic) is describing?
Role	Is the Source or Subject a Controller, Actuator, Controlled Process, or Sensor?
Behavior Type	For the given role, which type(s) of behavior does it exhibit? See the lists on the previous pages for each control role.
Context	Provide a justification for categorizing the text (or graphic) in the chosen manner.

Table 6. Database Version of Control Model

See Figure 9		Description
1.	Controller	Which controller is being described in the text?

35

2.	Actuator	What mechanism(s) does the control have in order to affect the process?
3.	Cntl'd Process	What process does the controller have control over?
4.	Sensor	What type of feedback does the controller receive about the process it controls?
5.	Process Model	What states and variables does the controller know about the process it controls?
6.	Cntl Algorithm	Does the controller use an algorithm or procedure to generate action?
7.	Control Actions	What types of action can the controller generate?
8.	Controller Status	Does the controller provide feedback to higher level controllers?
9.	Control Input	Does the controller receive set points or other types of commands?
10.	Controller Output	Does the controller have output other than through the actuator? This often includes transmission of information to other controllers.
11.	External Input	Does the controller receive external input, either in terms of other system information or other controller action(s), or other (e.g. a power source)?
12.	Alt Controller	Does the process receive action from controllers other than in item 1, 2?
13.	Process Input	Does the process require external input to function? Examples include pressure, power, and heat.
14.	Proc Disturbance	What environmental factors does the process interact with?
15.	Process Output	Does the system require that the process output something to other components? (e.g. power, pressure)

In order to obtain a "complete" model of the ConOps, this model development approach should be applied recursively over the entire ConOps document. The keywords, with associated questions and comments (Table 5 and Table 6), can be applied to individual sentences, paragraphs, and/or graphical depictions[7].

5.2.2 Formal Expression of Model Development

This section develops a mathematical formalism that is intended to achieve two ends. First, the formalism allows the analyst to achieve more rigor than using the text- and graphics-based

---

[7] For example, a graphical depiction of information flows in a ConOps could provide the information
  necessary for model development

descriptions in the previous section. That is, the formalism allows the analyst to apply the technique in a repeatable fashion and develop a model that is easier to query. Second, the formalism lends itself to tool development in future work. The model-based systems engineering paradigm is ultimately focused on developing both the theory and tools to assist in managing complexity and assisting in development and analysis. A control-theoretic formalism of a Concept of Operations is as follows.

Using the development from Mesarovic et al. [1970], the basic feature of a hierarchical system is that at one level a subsystem applies constraints to, and receives feedback from, a lower level subsystem (see Figure 6). These features are now re-formulated in control-theoretic terms. A controller, actuator, process, and sensor at some level i will be denoted as $C^i, K^i, P^i, C^i, K^i, P^i$, and L^i, L^i, respectively. What is important to note in hierarchy theory is that the *controlled process* at one level can actually be a controller at the level below. That is

$$\forall i \neq 1, \qquad \mathcal{P}^i = \mathcal{C}^{i-1} \tag{1}$$

where the controller $C^{i-1}C^{i-1}$ has its own actuators, controlled processes, and sensors $\left(K^{i-1}, P^{i-1}, L^{i-1}\right)\left(K^{i-1}, P^{i-1}, L^{i-1}\right)$. Without loss of generality, the following development drops the superscript notation, except where necessary.

The entire ConOps document is denoted, CC, and the document consists of a structured group of elements that contain information about how the concept should behave. These elements could be sentences, paragraphs, or graphical objects such as figures that depict information flows or sequences of events. These "information elements" are denoted $\mathbb{I}$. The ConOps is then the set:

$$\mathbb{C} = \left\{ \bigcup_i \mathcal{I}_i \,\middle|\, (\mathcal{I}_i \in \mathbb{I}) \land (\mathcal{I}_i \cap \mathcal{I}_i = \emptyset), \forall i \neq j; i, j \in N \right\} \tag{2}$$

where $\mathbb{I}$ is any element that conveys information about operational concepts in terms of prose or graphics, and N is the total number of elements in a document. Equation 5 ensures that the model generation process does not use duplicate or overlapping information. The process is repeated recursively over each element $\mathbb{I}$, resulting in completeness while avoiding duplication and potential inconsistency. Decomposing a document into a coherent, mutually exclusive set of information elements requires some basic understanding of grammar[8].

5.2.2.1   Identifying the Components of the Model

The following formalism provides guidance for how to identify the elements necessary to generate a model of the concept. Each information element I is defined as a tuple $(S, R, B, A)(S, R, B, A)$ where:

S S is the source or subject of the information object $\mathbb{I}$. Any complete sentence in English should have a subject and predicate. Identifying $\mathcal{S}$ is often similar to identifying the subject of a sentence in grammar, while identifying the predicate of the sentence yields the rest of the

---

[8] To a certain extent, decomposing a document also depends on the competency of the original authors.

information in the tuple. The information object, II, could be a set of sentences or paragraphs that share the same subject. Identifying the subject of a graphical object may not be as straightforward as in natural language text, and graphics may contain multiple subjects or sources of responsibility[9]. Much of the information of a graphical object can be inferred using the associated text that refers to the graphic or by using the model or data that underlies the graphics.

RR is the responsibility of the subject in control-theoretic terms.

$$\mathcal{R} \in \{\mathcal{C}, \mathcal{K}, \mathcal{P}, \mathcal{L}\} \tag{3}$$

where:

$$\mathcal{C} := \text{controller,}$$
$$\mathcal{K} := \text{actuator,}$$
$$\mathcal{P} := \text{process being controlled,}$$
$$\mathcal{L} := \text{sensor.}$$

BBis the type of behavior prescribed to the source, and $B_R B_R$ represents the possible behaviors ascribed to an arbitrary responsibility RR. The available behavior types are

$$\mathcal{B} \in \{\mathcal{B}_c, \mathcal{B}_\mathcal{K}, \mathcal{B}_\mathcal{P}, \mathcal{B}_\mathcal{L}\} \tag{4}$$

The controller $(B_c)(B_c)$ represents a transformation $F_c F_c$ from input signals, $I_c I_c$, to output signals, $O_c O_c$. The input-output model for the controller is

$$\mathcal{B}_c = \{(\mathcal{I}_c, \mathcal{O}_c) | \mathcal{O}_c = \mathcal{F}_c \times \mathcal{I}_c, \mathcal{F}_c = f(G, \rho)\} \tag{5}$$

where the set of controller inputs, $\mathcal{I}_c$, consists of feedback information, $\mathcal{I}_s$, communications from other controllers, $\mathcal{I}_o$, and higher level commands or set points, $\mathcal{R}$. The set of controller outputs, $\mathcal{O}_c$, consists of the available control actions, $\mathcal{O}_a$ and information transmission to other controllers, $\mathcal{O}_o$. The control function, transformation $\mathcal{F}_c$, is a function of the algorithm (or procedure, decision-making process, or policy), G, and process model, ρ. Furthermore, formation and maintenance of the process model is performed via feedback and external information sources, ρ_s and ρ_e, respectively. Therefore, in terms of operational concepts, a controller is comprised of a sub-set of behaviors related to processing inputs and generating outputs:

$$\mathcal{B}_c \in \left\{\mathcal{B}_{\mathcal{I}_c}, \mathcal{B}_{\mathcal{O}_c}\right\} \tag{6}$$

$$\mathcal{B}_{\mathcal{I}_c}: \mathcal{I}_c \to \rho,$$
$$\mathcal{B}_{\mathcal{O}_c}: \rho \times G \to \mathcal{O}_c.$$

---

[9] In fact, this is one of the benefits of using graphical depictions and of model-based systems engineering in general. That is, graphical depictions allow for the storage of many different kinds of information in one concise space.

Actuator behavior $(\mathcal{B}_{\mathcal{K}})$ is a signal mapping from controller commands, O_c, to manipulated process variables, $\mathcal{V}_c$. The behavioral model of the actuator is

$$\mathcal{B}_{\mathcal{K}} = \{v | O_c \to \mathcal{V}_c\} \tag{7}$$

Process behavior $(\mathcal{B}_p)$ is a transformation F p from input signals, I p , to output signals,

$$\mathbf{B_p} = \{(\mathcal{I}_p, O_p) | O_p = \mathcal{F}_p \times \mathcal{I}_p, \mathcal{F}_p = f(\mathcal{V})\} \tag{8}$$

where the set of process inputs, $\mathcal{I}_p$, consists of an actuator signal intended to manipulate certain variables of the process, $\mathcal{V}_m$, actions from other controllers, $\mathcal{V}_A$, external process inputs, $\mathcal{I}_e$, and disturbances, $\mathcal{D}$. $\mathcal{F}_p$ is the process dynamics of the system. The set of process outputs, O_p, consists of external outputs, O_e and signals related to variables under control, $\mathcal{V}_c$. A controlled process exhibits two behavior types:

$$\mathcal{B}_p \in \left\{ \mathcal{B}_{\mathcal{C}_p}, \mathcal{B}_{\mathcal{D}_p} \right\} \tag{9}$$

$$\mathcal{B}_{\mathcal{C}_c}: \mathcal{I}_c^i \to \mathcal{I}_c^{i-1},$$
$$\mathcal{B}_{\mathcal{D}_p}: \mathcal{V}_m \times \mathcal{V}_A \times \mathcal{I}_e \times \mathcal{D} \to O_e \times \mathcal{V}_c.$$

Note again that for any controller, C^i (re-introducing the superscript notation), the role of its controlled process is also that of a controller of a lower level process when $i \neq 1$. In other words, process behavior of type $\mathcal{B}_{\mathcal{C}_p}$ should automatically trigger a recursion, where a lower level analysis produces another set of actuators, processes, and sensors. That is,

$$\mathcal{B}_{\mathcal{C}_p} \Rightarrow \mathcal{P}^i = C^{i-1} \tag{10}$$

subject to the rule in equation (4). Rather than exhibiting the dynamic behavior typically associated with process control, the transformation performed by the process is equivalent to the mapping performed by a controller $(\mathcal{F}_p^i \equiv \mathcal{F}_c^{i-1})$. Alternatively, at the bottom level,

$$i = 1 \leftrightarrow \mathcal{B}_p \equiv \mathcal{B}_{\mathcal{D}_p}. \tag{11}$$

Thus, if the process is at the lowest level of the system hierarchy, the transformation $\mathcal{F}_p$ represents the system dynamics. In the controls literature this mapping is often represented in state space as a dynamic, feedback control system of the form,

$$\dot{x}(t) = f(x(t), u(t))$$
$$y(t) = g(x(t), u(t)) \tag{12}$$

where $x(t)$ represents the dynamics of the system subject to control input $u(t)$. The second equation, $y(t)$, represents the feedback terms. In most engineered systems, the hazardous states

that should be eliminated or mitigated can be approximated by a vector of continuous or semi-continuous functions, as in equation (15)[10].

Sensor behavior $(\mathcal{B_L})$ represents a signal mapping from measured variables, $\mathcal{V}_m$, to controller inputs, $\mathcal{I}_s$. The behavioral model of the actuator is

$$\mathcal{L} = \{i \mid \mathcal{V}_m \to \mathcal{I}_s\} \tag{13}$$

The type of feedback controller input defines the behavior of the sensor, which consist of the set

$$\mathcal{B_L} \in \left\{\mathcal{B}_{\mathcal{L}_c}, \mathcal{B}_{\mathcal{L}_d}, \mathcal{B}_{\mathcal{L}_s}\right\} \tag{14}$$

$$\mathcal{B}_{\mathcal{L}_c}: \mathcal{V}_m \to \mathcal{I}_{s,c}$$
$$\mathcal{B}_{\mathcal{L}_d}: \mathcal{V}_m \to \mathcal{I}_{s,d}$$
$$\mathcal{B}_{\mathcal{L}_s}: \mathcal{V}_m \to \mathcal{I}_{s,s}$$

where $\mathcal{I}_{s,c}$ is a continuous signal (or semi-continuous digital signal) representing the evolution of the process, $\mathcal{I}_{s,d}$ is discrete data representing a state transition of the process $\mathcal{P}$, and $\mathcal{I}_{s,s}$ is a synthesis of measured process variables into a lower-dimensional data stream $\left(\dim(\mathcal{I}_{s,s}) < \dim(\mathcal{V}_m)\right)$.

$\mathcal{A}$ is the context or set of assumptions, which provides the analysts' justification for assigning the first three elements to the quadruple $\mathcal{I}$. Context can simply be a textual rationale, reference to other parts of the ConOps, reference to other documents, or other means of inference. $\mathcal{A}$ is ultimately not a part of the control-theoretic model development, but it is an important aspect of making explicit the often undocumented assumptions that are present during concept development.

Each information element, $\mathcal{I}$, can easily be stored in a database consisting of the source or subject, its responsibility within the system, its behavior, and the assumptions or context used to define, identify, or classify the information element. Not only can the model that results from the aggregation of tuples be analyzed (see next section), but also the model is easily traceable back to the original ConOps.

Because this report advocates for a systems approach to developing a concept, it is insufficient to simply identify and store all the model information according to the tuple $(\mathcal{S}, \mathcal{R}, \mathcal{B}, \mathcal{A})$ and then analyze their individual behavior. The analyst must also identify the relationships between all of the elements and then analyze both component behavior and interactions between those components.

5.2.2.2   Synthesizing Information into Hierarchical System Model

The previous development involves parsing the ConOps by mapping information elements I into control-theoretic constituents defined by the tuple $(\mathcal{S}, \mathcal{R}, \mathcal{B}, \mathcal{A})$. The resulting tuples do not,

---

[10] In other systems that have very real hazards, such as investment banking, such dynamic equations may not exist. However, a mapping $\mathcal{F}_p$ for inputs to outputs should still exist.

however, by themselves represent a model of the entire concept. The above analysis should result in a set of controllers, each with its own actuators, processes, and sensors. Section 5.1 describes the importance of emergence and hierarchy in systems theory, but it is important to develop that theory further here. In systems theory it is inappropriate to analyze individual control loops and then make a determination about the overall behavior of the system. Furthermore, it is inappropriate to analyze individual components like sensors, actuators, or controllers.

Rather, the behavior of the system can only be determined in the context of all the components and their interactions. Instead of focusing solely on understanding the behavior of each component, the relevant issue here relates to how the individual control elements relate to each other. This section develops a formalism, based on hierarchy theory, to determine the relationships between control components.

Though the following section presents a formalism for checking the consistency across the hierarchy, this section presents heuristics for *identifying* or constructing the hierarchy based on information contained in the ConOps. This research proposes the use of several abstractions that can be used to determine the "vertical" and "horizontal" relationships between control components. This section also introduces the formal mathematical notation to be used later.

Several different (but related) notions of hierarchy, or abstraction, may be used to generate a system control model from the individual control loops generated from—and stored in—the tuples $(\mathcal{S}, \mathcal{R}, \mathcal{B}, \mathcal{A})$. For example, Mesarovic and Takahara [1975] describe three types of hierarchies: strata or levels of description, layers or levels of decision complexity, and echelons or organizational decomposition. In the controls literature the hierarchy is typically layered in terms of time scale, for example, scheduling (weeks), system-wide optimization (days); local optimization (hours); supervisory, predictive, or advanced control (minutes); and regulatory control (seconds) [Skogestad, 2004]. The echelon hierarchy—described by Mesarovic and used in the controls literature [e.g. Morari et al., 1980]—is often used to decompose a system using the notion of decision-making authority. That is, some decision-making units are influenced or controlled by others.

A specific characteristic of the echelon hierarchy is that there are many elements within a given level, which implies another dimension of organization. Intent Specifications [Leveson, 2000b] organize system information according to three types of hierarchy: level of intent, part-whole abstractions, and refinement. Part-whole abstraction provides another horizontal decomposition of the system. That is, while decision complexity, time scale, or authority defines a hierarchy *vertically*, part-whole abstractions describe the organization and relationships *horizontally* within a given level.

5.2.2.3 Hierarchical Control Structure

Any level of a system can be represented as $: X \rightarrow Y$, a mapping from a set of outside stimuli, X to a set of responses, Y .

$$X = X_1 \times X_2 \times \cdots \times X_n \text{, and} Y = Y_1 \times Y_2 \times \cdots \times Y_n \tag{15}$$

The i^{th} level of the system is then the mapping

$$S_i : X_i \rightarrow Y_i \tag{16}$$

The previous step of the analysis should have identified a set of stimuli and responses, in particular in the controller and process behaviors, $\mathcal{B}_c$ and $\mathcal{B}_p$, respectively. The i^{th} mapping is then comprised of the controller inputs and outputs. From equation (8), the mapping consists of:

$$X_i = \{\mathcal{I}_o^i, \mathcal{I}_S^i, \mathcal{R}^i\} \tag{17}$$

$$Y_i = \{\mathcal{O}_a^i, \mathcal{O}_o^i, \mathcal{F}^i\} \tag{18}$$

Building on the concept introduced in equation (4),

$$
\begin{array}{ll}
(i) & S_i\colon \mathcal{I}_o^i \times \mathcal{I}_S^i \to \mathcal{O}_a^i \times \mathcal{O}_o^i \ \text{if}\ i = n, \\
(ii) & S_i\colon \mathcal{I}_o^i \times \mathcal{I}_S^i \times \mathcal{R}^i \to \mathcal{O}_a^i \times \mathcal{O}_o^i \times \mathcal{F}^i \ \text{if}\ 1 < i < n, \\
(iii) & S_i\colon \mathcal{I}_o^i \times \mathcal{R}^i \to \mathcal{O}_o^i \times \mathcal{F}^i \ \text{if}\ i = 1.
\end{array} \tag{19}
$$

The set of systems S_i, $1 \le i \le n$, is a hierarchical control structure if there exist two families of mappings $h_i\colon Y_i \to \mathcal{F}_{i+1}$, $1 \le i \le n$ and $c_i\colon Y_i \to \mathcal{R}_{i-1}$, $1 \le i \le n$, such that for each x in X and $y = S(x)$:

$$
\begin{array}{ll}
(i) & y_n = S_n(x_n, h_{n-1}(y_{n-1})), \\
(ii) & y_i = S_i(x_i, c_{i+1}(y_{i+1}), h_{i-1}(y_{i-1}))\ 1 < i < n, \\
(iii) & y_1 = S_1(x_1, c_2(y_2)).
\end{array} \tag{20}
$$

5.2.2.4 Safety and Authority within Control Structure

The structure described in equations (19)–(23) represents a general description of input-output systems and their hierarchical relationships. Each sub-system, S_i, could simply contain a feedback control algorithm, for example PID[11]. However, in systems with sufficient complexity there often exists a set of decisions, and the control agent must select among alternatives. That is, the mapping $S_i\colon \mathcal{R}^i \to \mathcal{R}^{i-1}$ is a decision-making unit. In this type of hierarchy, there exists a family of decision problems $\mathcal{D}_i(Y_i)$, $Y_i \in \mathcal{R}^i$ and a transformation $\mathcal{K}_i$ such that for any input Y_i the output $Y_{i-1} = S_i(Y_i)$ is given as $Y_{i-1} = \mathcal{K}_i(x_i)$ where x_i is a solution of the decision problem $\mathcal{D}_i(Y_i)$. The inputs $Y_i \in \mathcal{R}^i$ from the decision-making unit immediately above act as a parameter in the decision problem of sub-system S_i. Alternatively, the outputs Y_{i-1} obtained from the transformation $\mathcal{K}_i$ are parameters for the lower level decision units. Such a decision-making hierarchy exists when S_i has "priority of action" or "control authority" over S_j.

In the management literature, the hierarchical structure of an organization is described in terms of general decision making. For example, a decision making problem under conditions of uncertainty may be specified as a satisfaction problem by the 4-tuple (g, τ, X^f, Ω): find a solution x in the feasible set X^f such that for all uncertainties ω in Ω:

---

[11] PID $\triangleq$ Proportional-Integral-Derivative, a form of feedback control widely and successfully used in industrial systems.

$$g(X, \omega) \le \tau(\omega) \tag{21}$$

where $\le$ is a given relation, τ is a tolerance function, and g is an objective function.

In safety-driven design, the organization need not be defined in terms of general decisions. Rather, the safety-driven approach focuses on the identification and prevention of *unsafe* decisions; in the paradigm of hierarchical control, these unsafe decisions are defined more precisely as unsafe control actions. Building upon the formalism of unsafe control developed by Thomas [2013], an unsafe control action in the STAMP accident model can be expressed formally as a 4-tuple $(\mathcal{C}^i, T, \mathfrak{C}, Co)$ where:

- $\mathcal{C}^i$ is the source controller that can issue control actions in the system. The controller may be automated or human.
- T is the type of control action. There are two possible types: Provided describes a control action that is issued by the controller while Not Provided describes a control action that is not issued.
- $\mathfrak{C}$ is the control action (i.e. command) that is output by the controller.
- Co is the context in which the control action is or is not provided.

Section 5.3 further develops this formalism; the tuple provides a rigorous way to trace controller actions to hazards and vice versa. For the purpose of merely developing the model, an important omission[12] from the above 4-tuple is the *destination* of the control action. Thus, the definition is extended to include a destination $\mathfrak{C}_i^j$, $\quad i, j \in N$, which represents a control action from the j^{th} source to the i^{th} destination.

Without loss of generality, a three-layer "decision" hierarchy is then:

$$S_1: \mathfrak{W}_1 \times \mathfrak{C}_1^3 \times \mathfrak{C}_1^2 \to \mathfrak{C}_1^1, \tag{22}$$

where $\mathfrak{W}_1$ is feedback information from the *controlled process*. $\mathfrak{C}_1^2$ and $\mathfrak{C}_1^3$ provide constraints on the set of actions available to controller $\mathcal{C}^1$. Depending on the architecture of the control system, this level might receive constraints only from the immediate level above it. That is, $\mathfrak{C}_i^j = \emptyset$ for $j - i \ge 2$ for completely stratified systems.

The second layer is represented by a mapping

$$S_2: \mathfrak{W}_2 \times \mathfrak{C}_2^3 \to \mathfrak{C}_1^2, \tag{23}$$

and $\mathfrak{W}_2$ is feedback available to the second-level controller. At this level, the controller may not have direct access to hazard variable states $V(\mathcal{C}^1)$ but may have additional information about the environment that the lower level controller does not have access to. The final layer of this three-layer hierarchy is

$$S_3: \mathfrak{W}_3 \times \mathfrak{C}_1^3 \times \mathfrak{C}_2^3, \tag{24}$$

---

[12] An omission that Thomas acknowledges, and can be extended.

At some point in the hierarchy, the system must have access to information about hazardous states. That is, W i should provide information about the variables relevant to system hazards present in the controlled process of controller $\mathcal{C}^1$.

Consider a very simplified example in aircraft guidance and navigation. The flight crew (S_3) is informed of convective weather $(\mathfrak{W}_3)$ and inputs a new series of waypoints into the Flight Management System (S_2). The Flight Management System then sends commands to the aileron hydraulics (S_1). The local aileron control system uses sensor input about its position $(\mathfrak{W}_1)$ to adjust pressure in the hydraulics. The Flight Management System uses aircraft position data $(\mathfrak{W}_2)$ to update commands to the local aileron controller as well as to send position information to the Flight Crew. In this case, $\mathfrak{W}_2 \subset \mathfrak{W}_1$.

In the above aircraft example it is apparent that many more components[13] are required to direct the aircraft in the proper direction. These lower-level control components are highly coupled, and the guidance of an aircraft depends on the simultaneous manipulation of many variables $\mathcal{V}_m$. This relationship suggests another dimension of decomposition that must be accounted for when identifying and synthesizing a control structure model.

5.2.2.5   Other Vertical Relationships

The level of a given controller may not be obvious at this early stage of concept development. That is, decision-making priority may not be evident, or may not have even been defined yet, in an early concept of operations. Other notions of vertical decomposition are described here in order to guide this process of identifying the appropriate level of a control agent.

Decision Complexity represents another type of decomposition in control hierarchies. Increasingly complex decisions tend to lack well-defined and complete specification of uncertainties, input conditions, problem constraints, and processes involved in transforming input conditions into desired output states [Simon, 1977]. Such complex decisions often require selection among multiple alternatives and involve interrelated factors with time dependence and nonlinearities such as feedback lag, delayed effects, singularities, tipping points [Forrester, 1987; Sterman, 1994]. Decision complexity in systems theory often implies that a decision at one level involves processing, understanding, and coordinating decisions at a lower level (see equation 23).

Time scale constitutes another form of vertical decomposition. For example, an aileron control (sub) system in an aircraft measures and adjusts commands on the order of fractions of seconds; a pilot (in a highly automated civilian aircraft) might be adjusting flight plans or profiles on the order of minutes or hours. In this case there is already a built-in priority of action in most aircraft systems, but these timing differences serve as another indication of a hierarchical decomposition. A heuristic for vertical decomposition of control agents is then

$$\left\{ S_i = f(t_i), S_j = f\left(t_j\right) \middle| t_i \gg t_j \right\} \Rightarrow i > j \tag{25}$$

---

[13] Ailerons, elevators, rudders, spoilers, thrust, and many others are required to control an aircraft trajectory.

where $S = f(t)$ represents a control output as a function of time interval, t. The following notions describe additional aspects of control hierarchy.

(1) Higher level units are concerned with larger portions of the system, which can be modeled using the STAMP notion of a process model, along with the concepts of aggregation and set theory.

$$\{ \exists \rho^i \in \mathcal{C}^i, \rho^j \in \mathcal{C}^j \,|\, (\rho^i \supset \rho^j) \vee (\rho^j \in \rho^i) \} \Rightarrow i > j \tag{26}$$

That is, if a process model of one controller (i) either is a superset of, or contains, the process model of another controller (j) then controller i is supremal to controller j. If aspects of ρ^i are in ρ^j and vice versa, but the superset condition does not hold, then the relationship is most likely horizontal and not vertical. The next sub-section describes horizontal decomposition.

(2) Higher level units are not only concerned with the slower aspects of the systems' operation (see equation 28) but also: the exchange with the environment takes place at a lower frequency, the dynamics of concern is slower, and the period between decision time is longer.

$$\left. \begin{array}{c} \lambda_i(f(x_i)) \ll \lambda_j\left(f\left(x_j\right)\right) \\ \lambda_i(\mathcal{D}_i) \ll \lambda_j\left(\mathcal{D}_j\right) \\ T(f(x_i)) \ll T(\mathcal{C}^j) \end{array} \right\} \Rightarrow i > j \tag{27}$$

where $\lambda_k(f(x_k))$ is a fundamental mode of the dynamics of subsystem S_k, $\lambda_k(\mathcal{D}_k)$ is a fundamental mode of the disturbances to subsystem S_k, and $T(\mathcal{C}_k)$ is the time between control actions associated with S_k.

(3) *Abstraction hierarchies* [Rasmussen, 1986] decompose the system in terms of level of description. A control agent may be concerned with functional purpose, abstract function, generalized function, physical function, or physical form. The Functional Purpose level describes the goals and purposes of the system, and systems typically include more than one system goal such that the goals conflict or complement each other. The relationships between the goals indicate potential trade-offs and constraints within the work domain of the system. For example, the goals of a flight planner might be to achieve a desired route while trading off between flight time versus fuel consumption. The Abstract Function level describes the underlying laws and principles that govern the goals of the system. These are typically empirical or theoretical laws in an engineered system, but economic or judicial principles underlie a social system. Aircraft flight is governed by laws related to thrust, lift, and drag.

The Generalized Function level explains the processes involved in the laws and principles found at the Abstract Function level, i.e. how each abstract function is achieved. Causal relationships exist between the elements found at the

45

Generalized Function level. To generate thrust, a turbofan uses fuel injection and intake air, which has implications at the Functional Purpose level. The Physical Function level reveals the physical components or equipment associated with the processes identified at the Generalized Function level. The capabilities and limitations of the components such as maximum capacity have implications all the way up the hierarchy. The Physical Form level describes the condition, location, and physical appearance of the components shown at the PFn level. In the aircraft example, the wings, turbofans, and fuselage are arranged in a specific manner, basically illustrating the location of the components.

5.2.2.6   Horizontal Decomposition

Developing a systems-theoretic model of a concept lies primarily in identifying the vertical relationships described above. However, within any level of control, S_i, there may exist a number of individual control agents, controlled processes, and other entities. The starting point is to recognize vertical position of the units according decision-making priority, abstraction, or other types of vertical relationships. Decomposition within one level of a hierarchy can then be done in terms of part-whole abstraction [Rasmussen, 1986; Leveson, 2000b] or echelons [Mesarovic et al., 1970].

Part-whole abstractions involve refinement and its opposite, aggregation. An intuitive description of aggregation is as follows. Suppose that P_1 is a mathematical description of a physical system using a given set of variables, and P_2 is a consistent description of the same system using a smaller set of variables. Then P_2 is termed an aggregate model for P_1, and the variables of the system P_2, are termed aggregate variables. Any of the variables within the refined model P_2 can then be said to have horizontal relationships in the control hierarchy. The following section explores some consistency properties related to aggregation and horizontal decomposition.

Another way to reason about horizontal relationships is to consider span of control. In an organizational hierarchy (also called a "Multi-echelon hierarchy"), any agents under the same span of control will have horizontal relationships. Building on the previous formalization of hierarchical control structures, and using the notation of Mesarovic and Takahara [1975], the following development formalizes the horizontal relationships among control agents. If C is a (finite) family of sub-systems $S_i, i \in N$, where N is a finite set, and if $>$ is a strict partial ordering of N, then $(C, >)$ is a hierarchy of systems. If $(C, >)$ is a hierarchy of control systems, and the ordering $>$ is such that $i > j$ iff S_i has priority of action over S_j, then $(C, >)$ is a control structure hierarchy.

Echelons in a control structure hierarchy $(C, >)$ are recognized in terms of the ordering $>$, representing priority of action. The first echelon units are the minimal units of C; the family

$$C^1 = \{ S_i \mid i \in N_1 \}$$

is the first echelon, where

$$N_\downarrow 1 = \{ i \mid i \text{ "is a minimal element of" } N \}.$$

The i^{th} echelon units are the minimal units of C when all lower echelons are omitted; the family

$$C^\uparrow i = \{ S_\downarrow k \mid k \in N_\downarrow i \}$$

is the i^{th} echelon, where

$$N_i = \{k| \; k \text{ is a minimal element of } N - [N_1 \cup N_2 \cup \; ... \; \cup N_{i-1}]\}.$$

Finally, define multi-echelon control structures as a subclass of general control hierarchies. A hierarchy of control systems $(C, >)$ is a multi-echelon control structure if, $\forall i, j \in N$, there is at most a unique $k \in N$ such that $\forall l \in N,$

$$l > i \text{ and } l > j \; \Rightarrow \; l > k.$$

This condition implies that any member of C has at most one unit of the immediately higher echelon which has control authority over it. In other words, a strict multi-echelon control structure is a "pyramid" structure, which is rare in real systems. That is, there is almost always overlap in control authority, often for good reason. In fact, there is often a trade-off between the simplicity of having the strict multi-echelon property and the availability of multiple controllers. The following section describes methods for assigning control responsibility and for identifying issues with coordination and consistency among multiple controllers. Regardless of this trade-off, systematically identifying the span of control provides another way of identifying horizontal relationships. If multiple entities (control agents) respond to the action of the same C^i (that is, they are under the span of control of the same controller) then they are on the same level of the hierarchy and have horizontal relationships.

Continuing the aircraft example, the FMS tries to coordinate between many processes in order to achieve some flight path objective. These lower-level processes and associated control systems are dynamically coupled and operate in parallel, for example adjusting thrust while manipulating (multiple) control surfaces to achieve a smooth turn or climb. In the case of an aircraft, the control structure is not a strict multi-echelon hierarchy, because lower level control systems must respond to either FMS signals, manual pilot inputs, or both.

5.3   Systems-Theoretic Analysis of Model

Much of the control- and systems-theoretical foundation used to develop the model of the concept is also used to guide the analysis of the model. However, the focus shifts from modeling to identifying potential causal factors and invalid assumptions. According to Leveson [2012, p.81], there are several fundamental vulnerabilities in a hierarchical system:

> At each level of the hierarchical control structure, inadequate control may result from missing constraints (unassigned responsibility for safety), inadequate safety control commands, commands that were not executed correctly at a lower level, or inadequately communicated or processed feedback about constraint enforcement.

The control-theoretic approach emphasizes the importance of process models in enforcing adequate control: a process model must contain "the required relationship among the system variables (the control laws), the current state (the current values of the system variables), and the ways the process can change state" [Leveson, 2012, p.87], or the dynamics of the process. The four fundamental requirements of process control (see 1.a-d below) described in the previous sub-section must also be satisfied.

Once the control model of the ConOps has been built (previous sub-section), the hazardous scenarios and causal factors can be identified using these systems-theoretic views of accident causality. Specifically, the analysts, engineers, and stakeholders should ask:

(1) Are the control loops complete? That is, does each control loop satisfy a Goal Condition, Action Condition, Model Condition, and Observability Condition?

    (a) Goal Condition – what are the goal conditions? How can the goals violate safety constraints and safety responsibilities?

    (b) Action Condition – how does the controller affect the state of the system? Are the actuators adequate or appropriate given the process dynamics?

    (c) Model Condition – what states of the process must the controller ascertain? How are those states related or coupled dynamically? How does the process evolve?

    (d) Observability Condition – how does the controller ascertain the state of the system? Are the sensors adequate or appropriate given the process dynamics?

(2) Are the system-level safety responsibilities accounted for?

(3) Do control agent responsibilities conflict with safety responsibilities?

(4) Do multiple control agents have the same safety responsibility(ies)?

(5) Do multiple control agents have or require process model(s) of the same process(es)?

(6) Is a control agent responsible for multiple processes? If so, how are the process dynamics (de)coupled?

As in the previous section, these questions can be formalized, and further description is provided with the following formalization. Question 1 relates to completeness of the individual control loops, questions 2-3 relate to assigning safety-related responsibilities to various control agents, and questions 4-6 relate to coordination of multiple control agents. The analysis therefore proceeds through three basic areas, which are explored in the following sub-sections and depicted in the bottom left of Figure 10.

GENERAL, SYSTEMS-THEORETIC CONOPS ANALYSIS

SAFETY-DRIVEN DESIGN

ConOps

Identify System Hazards

Derive System Safety Constraints

Identify Control Concepts

Derive Refined Safety Constraints

Identify Hazardous Scenarios and Causal Factors

- Completeness Criteria: Sec 5.3.1
- Analyze Safety Resp: Sec 6.4.2
- Coordination & Consistency: Sec 5.3.3

Refine, Modify Control Structure

Figure 10. Proposed Methodology - Analysis

5.3.1 Completeness Criteria for Individual Control Loops

Completeness criteria for process control systems have been developed elsewhere [e.g. Leveson, 2000a]. While existing specification languages[14] are formal and executable, it is often not desirable and perhaps not possible, to specify an entire system formally. The purpose of the formalism here, during concept development, is not necessarily intended to support simulation but rather to provide a rigorous means for identifying gaps in the control loops. Because this report is intended to support safety-driven design and development, completeness criteria are explicitly linked to system hazards.

---

[14] There are many examples, including SpecTRM-RL [Leveson, 2000a], Statecharts [Harel, 1987], Prototype Verification System [Owre et al., 1996], AsmL [Barnett and Schulte, 2001], etc.

In addition to the original descriptions in Figure 9, the control loop in Figure 11 assigns the variables and mappings presented in the formalism of the previous section.
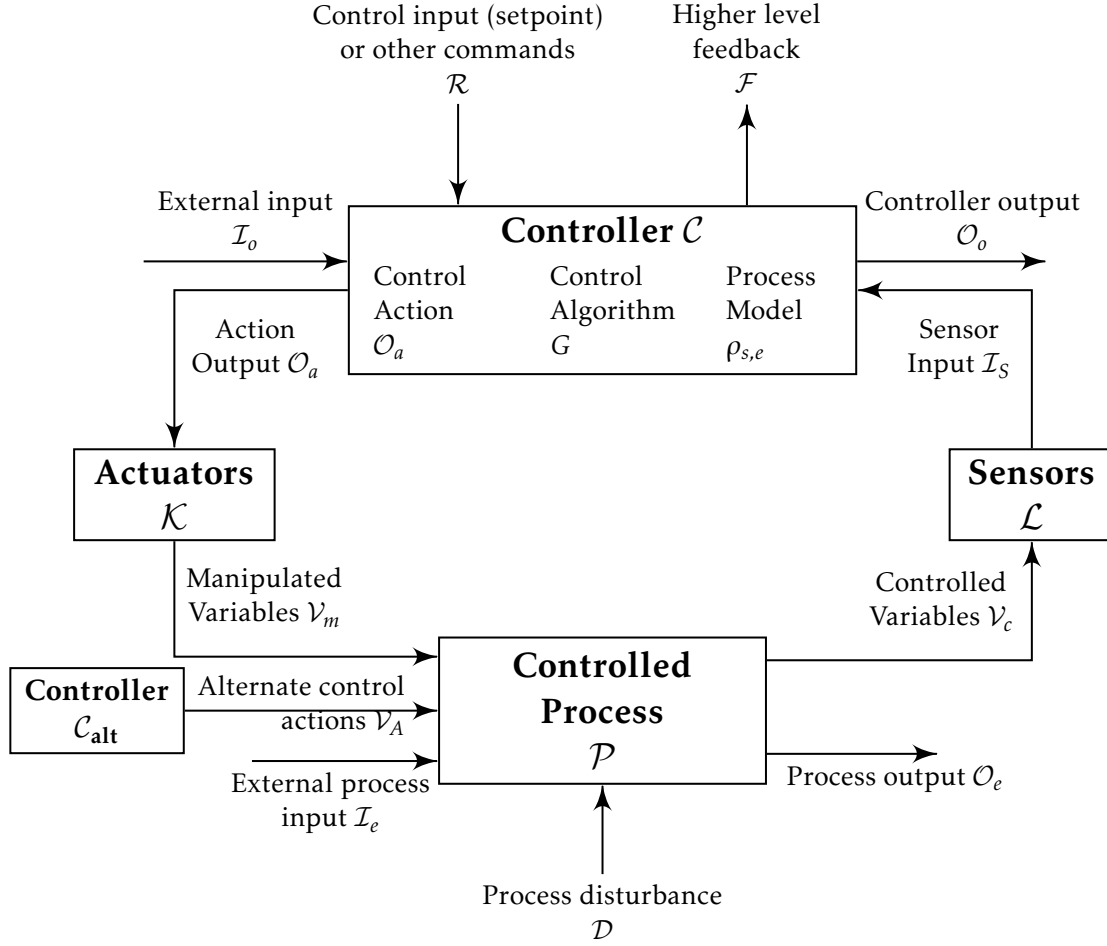
Figure 11. Generic Process Control Loop

Every system has a set of hazards, which are undesirable states of the system[15]. Define the set of states $V \subseteq X$, where X is the entire set of system state variables and V is a set of variables associated with hazard $\mathcal{H}$. For example, loss of separation (LOS) between aircraft occurs when they violate some minimum separation distance. Thus, $\mathcal{H}_{LOS}$ consists of the states, $V = \{x_1, y_1, h_1, x_2, y_2, h_2\}$, which are the current position of aircraft₁ and aircraft₂ in three dimensions each.

The control-theoretic approach to safety also assumes that processes are dynamic and can evolve over time or abruptly change. State dynamics typically take the form $\dot{X} = f(X, U, D)$ or $X_{k+1} = f(X_k, U_k, D_k)$. Recall that, for higher level control agents, the controlled process itself

---

[15] Section 5.4 has a more complete definition of a hazard, but the above definition serves for the current development.

50

could be a control agent and these dynamics may not be continuous. In the following formalism, process dynamics take the form

$$\mathcal{X} \times \mathcal{U} \times \mathcal{D} \mapsto \mathcal{X} .$$

(1) *Goal Condition*—the goal condition relates to preventing, and recovering from, hazardous states. In safety-driven design, the goal condition should seek, in order of decreasing priority, to (1) eliminate hazards, (2) avoid hazards, and (3) recover from hazards. That is, $\forall \mathcal{V} \in \mathcal{H}, \exists G, \mathcal{O}_a [G \rightarrow \mathcal{O}_a \wedge \mathcal{O}_a \Rightarrow \neg \mathcal{V}]$ and the negation of $\mathcal{V}$ can take several forms. Hazard elimination implies that the hazard does not exist $(\neg \mathcal{V} \Rightarrow \mathcal{H} = \emptyset)$, hazard reduction or avoidance implies that the system state will never become a hazardous state $(\neg \mathcal{V} \Rightarrow \mathcal{O}_{\downarrow} a \times x \mapsto x' \mid x, x^{\uparrow'} \notin \mathcal{H})$, and hazard control or recovery seeks to minimize the amount of exposure to the hazard $(\forall x \in \mathcal{H}, \exists \mathcal{O}_a [\mathcal{O}_a \times x \mapsto x', x' \notin \mathcal{H}],$ that minimizes the time from $x \rightarrow x')$.

(2) *Action Condition*—for every hazardous state variable, there is a signal that can manipulate that state, causing it to change. Furthermore, the actuator maps the controller output into the signal that can manipulate (hazardous) process states.

$$\forall \mathcal{V} \in \mathcal{H}, \exists \mathcal{K} [\mathcal{O}_a \times \mathcal{K} \mapsto \mathcal{V}_{\mathcal{K}} \wedge \mathcal{V}_{\mathcal{K}} \times \mathcal{V} \mapsto \mathcal{V}'] \tag{28}$$

where $\mathcal{X}'$ represents some change in state from $\mathcal{X}$.

(3) *Model Condition*—the model condition asserts that the control agent must have a model for how the system evolves $= [\mathcal{I}, \mathcal{V} \mid \mathcal{I} \times \mathcal{V} \mapsto \mathcal{V}']$, where $\mathcal{I} = \{\mathcal{C}, \mathcal{D}, \mathcal{E}\}$ are the control inputs, environmental disturbances, and other inputs from within the system, respectively. $\mathcal{V}'$ represents the evolution of the potentially hazardous state variables due to inputs and internal dynamics. Rather than consider every state in the system, in safety driven design emphasis is given to $\mathcal{V} \in \mathcal{H}$, the evolution of process variables that may lead to a hazard.

(4) *Observability Condition*— there is a signal that measures the states v, and the signal can also help discern the evolution from V to V ' due to action A
$\forall \mathcal{V} \in \mathcal{H}, \exists \mathcal{L} [\mathcal{O}_a \times \mathcal{V} \times \mathcal{L} \mapsto \mathcal{I}_{\mathcal{L}} \vee \mathcal{V} \mapsto \mathcal{V}' \Rightarrow \mathcal{I}_{\mathcal{L}} \times \rho \mapsto \rho'],$
where ρ' represents some update of the model ρ. The above definition of observability not only asserts that the signal must be updated for every change in the (hazardous) process states, the signal must also explicitly discern the change (or lack of change) due to action, $\mathcal{O}_a$.

5.3.2 Analyzing Safety-Related Responsibilities

This part of the analysis is intended to ensure that all hazards and safety constraints are accounted for in the hierarchical control structure and to identify goals and responsibilities that conflict with safety constraints.

Given a set of system hazards, $H_i \in \mathcal{H}$, then a safety constraint represents control over system behavior that prevents the hazardous states from occurring. That is,

$$\exists \sigma_i \in \Sigma, \sigma_i \Rightarrow \neg H_i \tag{29}$$

51

where Σ is a set of constraints, and σ_i is a safety constraint. There are two fundamental hazardous scenarios associated with system hazards and safety constraints. The first is that a safety constraint is unaccounted for. The second basic type of scenario is when the enforcement of a particular safety constraint (or general system goal) can cause a different hazard. In the framework of a hierarchical control structure, safety constraints and system goals are enforced via control actions, $\mathbb{C}$.

Let $\mathbf{P}(x, \mathbb{S})$ be defined for all pairs $(x, \mathbb{S})$, where

$$\mathbf{P}(x, \mathbb{S}) \equiv x \text{ results in} \mathbb{S} \tag{30}$$

The predicate $\mathbf{P}(x, \mathbb{S})$ is true, iff $\mathbb{S}$ is a defined condition or system state, and x is an action resulting in that condition.

A control structure that enforces every safety constraint is one that has available actions that result in equation (32). That is,

$$\forall \sigma_i \in \Sigma, \exists c \in \mathbb{C}[\mathbf{P}(c, \sigma_i)], \tag{31}$$

which states that, for every safety constraint there exists at least one available control action that causes the constraint to be true. A control structure with gaps is defined as a system that does not satisfy equation (34) and does not have appropriate control actions available to the various control agents.

Alternatively, there may be actions that conflict with safety constraints and actually cause the undesired hazard.

$$(\forall \mathbf{H_i} \in \mathcal{H})(\neg \exists c \in \mathbb{C})[\mathbb{P}(c, \mathbf{H_i}) \wedge \mathbb{P}(c, \mathcal{G})] \tag{32}$$

where $\mathcal{G}$ is a system goal state or system hazard such that $\mathcal{G} = H_j, i \neq j$. A system that does satisfy equation (35) is not necessarily a bad design or inherently unsafe. Rather, the identification of gaps (systems that do not satisfy equation 34) and conflicts (systems that do not satisfy equation 35) is intended to simply flag potential hazardous scenarios and causal factors.

These scenarios identify areas in the concept where architectural decisions, future design decisions, and refinement of safety-related and non-safety related requirements should be considered with great care. In fact, this analysis brings to bear design decisions and requirements that previously have not been identified as safety-related, as the example in Section 6 demonstrates. Thomas [2013] has developed a formal definition for hazardous control actions and a means for identifying them using STPA. Refer to Appendix C Formalism of Hazardous Control Actions for the formal definition of unsafe control actions and hazards. The preceding development assumes that hazardous control actions, which represent violations of safety constraints, can be identified using STPA. Identifying hazardous control actions becomes increasingly powerful as more design detail becomes available.

5.3.3 Coordination and Consistency

It is important to ensure that all safety responsibilities are accounted for, and often the system design results in either (a) one entity being responsible for enforcing multiple safety constraints or (b) multiple entities being responsible for enforcing the same safety constraints. Combinations of both do exist, especially in sufficiently complex systems.

Much theoretical work in systems theory has been dedicated to hierarchical control strategies related to problem (a) above. That is, much of the work attempts to ensure that there exists some policy guaranteeing that lower level control agents will achieve their individual objectives simultaneously with the higher level objectives [Mesarovic and Takahara, 1975]. The process control literature and other fields related to control theory have developed both the practice and theory of decomposing systems so that some coordination principles exist between the relatively de-coupled processes (e.g. [Acar and Ozguner, 1989]; [Morari and Stephanopoulos, 1980]; [Zheng et al., 1999]; [Skogestad, 2004]; [Tatjewski, 2008]).

This work neglects the latter problem—problem (b) above—where control decisions come from multiple sources and actuate on the same process variables. Often in process and chemical control, control structure designers are able to decouple the system sufficiently so that there is no overlap in responsibility. This is often not possible or not desirable in many complex socio-technical systems and/or systems with a high degree of dynamic coupling.

Safety-driven design is concerned not only with ensuring that coordination principles exist within the control structure but also in coordinating scenarios where multiple control agents have responsibility over the same process(es). Cowlagi and Saleh [2013] have suggested an approach and research direction for hazard analysis that builds upon some of the systems-theoretic concepts developed by Mesarovic [1970], as well as the approach to hazard analysis used here and originally proposed by Leveson [2004]. This report builds on that work but also includes the so-called "Multiple Controller" problem [Ishimatsu et al., 2010] in a formal way.

The first principle of safety-driven design relates to coordination of multiple controllers, which asserts that there must be some priority of action, or "leader", if processes (or control agents) can be manipulated by more than one source.

Define the priority of action function $\mathfrak{P}(\cdot, \cdot)$, as follows:

$$\mathfrak{P}(c, d) \Leftrightarrow c > d, \tag{33}$$

where the inequality implies that the action of c takes priority over action d. Conversely, if d takes priority over c or if there is no priority, then $\mathfrak{P}$ evaluates to false. That is

$$\neg \mathfrak{P}(c, d) \Leftrightarrow d > c \vee c = d. \tag{34}$$

Priority of action must also satisfy a transitive property,

$$\mathfrak{P}(a, b), \mathfrak{P}(b, c) \Rightarrow \mathfrak{P}(a, c). \tag{35}$$

Now define the action process predicate, $A(c, \mathcal{V})$ iff $c \times \mathcal{V} \to \mathcal{V}'$. That is, A is true whenever the action c can cause a change in state(s) from $\mathcal{V}$ to $\mathcal{V}'$. The coordination principle for two controllers (without loss of generality as long as transitive closure holds) is then,

$$(\forall c \in \mathfrak{C}_i)\left(\forall d \in \mathfrak{C}_j\right) \exists (\mathfrak{P}(c,d) \vee \mathfrak{P}(d,c))[A(c, \mathcal{V}_p) \wedge A(d, \mathcal{V}_p)] \tag{36}$$

where $\mathcal{V}_p$ is some system state, and actions c and d are generated from the i and j control agents, respectively. Equation (39) also enforces a control hierarchy. While the development and modeling effort may have identified control agents within the same level of the hierarchy, priority of action must be given if these control agents have overlapping control authority.

The second principle of safety-driven design is related to consistency of action. Any two controllers with safety-responsibilities related to the same process variables must ensure certain consistency characteristics. These controllers must have a consistent understanding or model of the current state, other actions that can affect that state, and how the state will evolve from those actions. The consistency principle for two controllers is:

$$\left(\forall v \in \mathcal{V}, \forall c \in \mathfrak{C}_i, \forall d \in \mathfrak{C}_j \,\middle|\, A(c, v) \wedge A(d, v)\right)\left[\rho_i(a, v) \equiv \rho_j(a, v) \wedge G_i \equiv G_j\right] \tag{37}$$

where $\rho_k(a, v)$ is the k^{th} control agent's model of state dynamics, v, subject to inputs, $a = \{c, d, \varepsilon\}$. The latter aspect of this set, ε, represents "other" inputs to the process. See equations (11) and (12) elsewhere for more detailed description of process inputs, which include but are not limited to environmental disturbances. G_k defines the goal condition of the k^{th} controller. The equivalence properties in equation (40) can be described using a property of aggregation called *dynamic exactness*. Suppose that the model ρ_1 is described by the state equation

$$\dot{x}(t) = Ax(t) + Bu(t) \tag{38}$$

and the model ρ_2 is described by

$$\dot{\hat{x}}(t) = M\hat{x}(t) + Nu(t) \tag{39}$$

In order for ρ_2 to be an aggregated model of ρ_1, it is required that

$$\hat{x}(t) = Cx(t) \tag{40}$$

for all t. This requirement is termed dynamic exactness. It is easy to see that dynamic exactness is achieved iff the matrix equations,

$$MC = CA, \tag{41}$$

54

$$N = CB. \tag{42}$$

are satisfied.

The equivalence property in equation (40) is relatively straightforward for many computer systems, where input-output behavior can be simulated and all relevant inputs are known. For human control agents, assuring—or even understanding—the consistency of mental models is more difficult, as is anticipating all possible input sequences for a complex computer system. As will be demonstrated in the following section, the coordination and consistency principles are intended to help the analysts identify potential scenarios that could arise due to independent sources of information, independent mathematical models, different goals or policies, and other factors. Such characteristics could be qualitative (such as for human operators or autonomous systems which are immature in terms of design detail) but could become increasingly quantitative (such as black box input-output models) as the concept matures.

As was the case with Analyzing Safety-Related Responsibilities (Section 5.3.2), the system is not necessarily unsafe if conditions in equation (39) and (40) do not hold. Rather, these conditions represent potentially hazardous aspects of a ConOps that require further inquiry and may need to be refined or modified in order to avoid potential conflicts.

The above description and formalism makes no distinction between deterministic or probabilistic signals or information, for example cases where a control or feedback signal only has some probability of reaching its intended destination. Because this analysis focuses on worst-case conditions and assumptions (see definitions in Table 7), the analyst must reason about system behavior if/when the signal does not reach its destination. Such reasoning is with respect to (a) control conditions, (b) fulfillment of safety-related responsibilities, and (c) coordination and consistency of control agents.

The framework presented in this section is based on control- and systems theory. It allows an analyst to systematically, rigorously interpret and decipher a natural-language description of a concept using guidewords and a series of generic roles and questions. Then, using a systems-theoretic view of accident causality, the framework asks a series of questions about the model to ensure completeness and consistency and to identify areas where further investigation is necessary. Results using this approach show that it identifies many more types of scenarios and factors than traditional PHA approaches (see Sections 6 and 8). Additionally, the approach helps analysts, engineers, and other stakeholders to identify and document more explicit and implicit assumptions about the concept.

5.4   *Using STECA in Early Systems Engineering*

The early phases of systems engineering involve identifying system objectives and criteria, defining top-level requirements, defining a system-level architecture, and then performing trade studies that ultimately lead to a design (see, for example, [Leveson, 2000b; de Weck et al., 2011; Kapurch, 2010; INCOSE, 2011]). As Section 2 argued, safety engineering should be integrated into these activities, and thus what follows is a brief explanation of those specific safety-related activities and their resulting artifacts. Table 7 depicts the relationships between safety-driven design activities and their counterparts in general systems engineering.

Table 7. General Systems Engineering and Safety-driven Design

General Systems Engineering	⇔	Safety-Driven Design
Identify System Objectives, Criteria	⇔	Identify Accidents and Hazards
Define Requirements	⇔	Define Safety Constraints
Define a System Architecture	⇔	Define a Hierarchical Safety Control Structure

Once the model and scenarios are identified, this information can be used to guide the system design. The methodology presented in Section 5.2 results in the definition of a control structure. This control structure should be representative of the descriptions contained in the ConOps, and this control structure could help define part of a baseline architecture for the system.

However, the analysis in Section 5.3 will identify potentially hazardous scenarios and associated causal factors. These scenarios drive the development of safety-related requirements and constraints (as well as functional requirements). The hazardous scenarios will also result in a refined or modified control structure. Alternative control structures should attempt to eliminate or mitigate against the hazardous scenarios found using the techniques and theory in Section 5.3.

Figure 12 shows the basic aspects of systems-theoretic early concept analysis. The figure captures the fact that the control model generated from the ConOps, and the subsequent analysis, form the inputs to STECA.

What is particularly noteworthy about the process is that modifying the control structure may fundamentally change some of the assumptions about the ConOps. For this reason, there is an upward arrow shown in Figure 12. This arrow depicts the fact that safety-driven design is, like any design effort, much more iterative and nonlinear. Changing the control structure should eliminate or mitigate certain hazards, but it may introduce new hazards. The figure also references the relevant sections later in this report.

GENERAL,
SYSTEMS-THEORETIC
CONOPS ANALYSIS
SAFETY-DRIVEN DESIGN

ConOps

Identify System Hazards

Derive System Safety Constraints

Identify Control Concepts

Derive Refined Safety Constraints
- Demonstration: Section 7.1

Identify Hazardous Scenarios and Causal Factors

Refine, Modify Control Structure
- Demonstration: Section 7.1.4
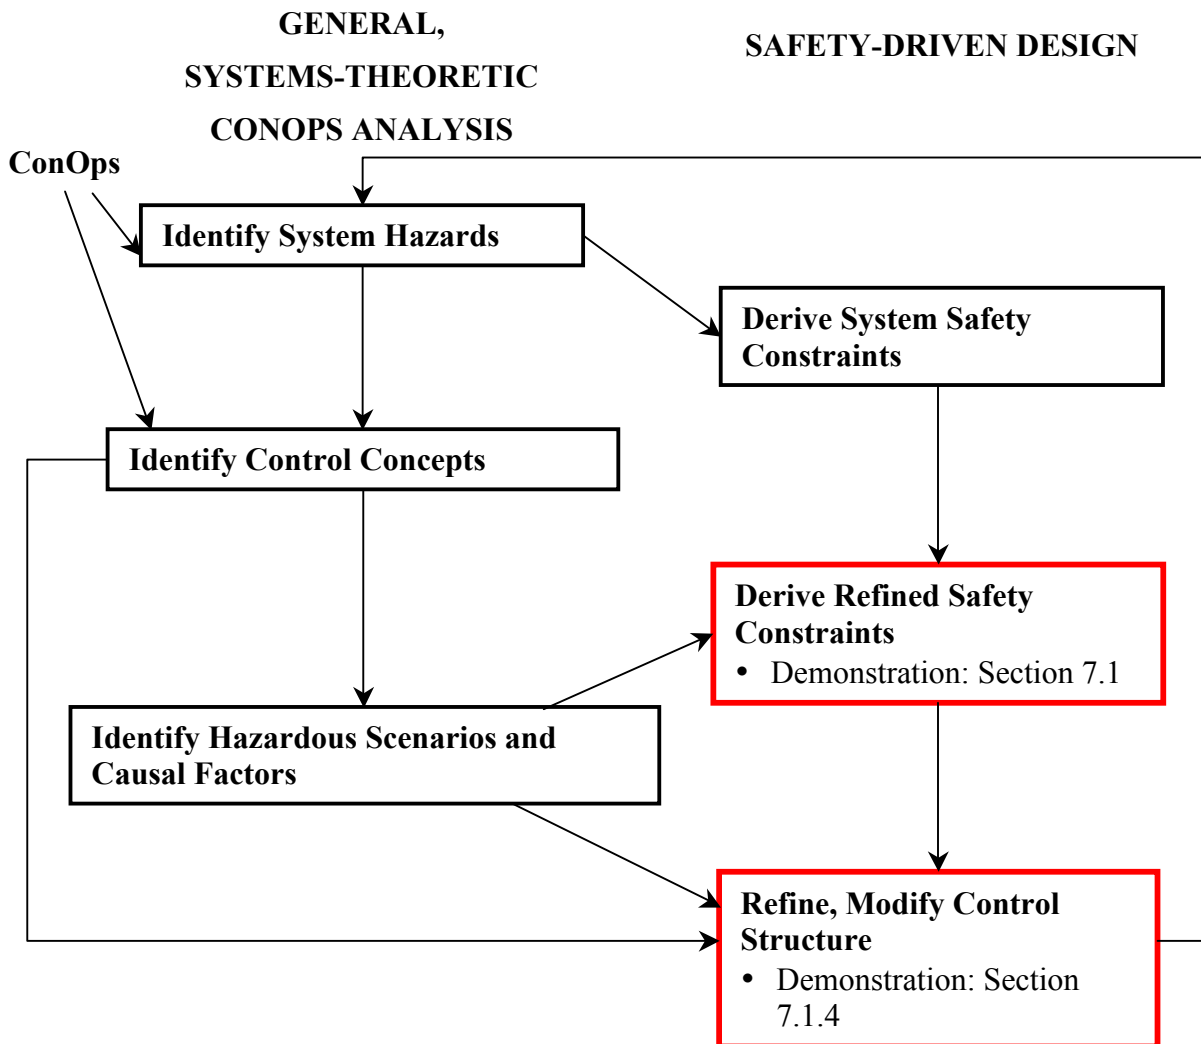
Figure 12. Proposed Methodology - STECA

The two processes in the bottom right of Figure 12, then, form the basis of safety-driven design: (1) derivation of refined safety constraints and requirements and (2) generation of a control structure and potential alternatives. Again, the control-theoretic modeling effort and subsequent analysis serve as inputs to these two activities.

Table 8 contains the definitions used throughout the rest of the report.

Table 8. Definition of Terms in Safety-Driven Design

Term	Definition
Accidents	undesired event that results in a loss, including human life or injury, damage to property, environmental pollution, etc.
Hazards	a system state or set of conditions that, together with a particular set of worst-case environmental conditions, will lead to an accident (loss)

Safety Constraint	limitation or restriction on system behavior. In safety-driven design the constraint should prevent hazardous behavior from occurring
Scenario	set of events, actions, or behaviors that can lead to a hazardous state (includes requirements flaws, inappropriate human behavior, dysfunctional interactions among components, and component failures)
Causal Factors	lower level or refined events, actions, or behaviors that contribute to a scenario
Control Structure	hierarchical structure where each level imposes constraints on the activity of the level beneath it
STECA	systems-theoretic early concept analysis, an analytical tool based on the STAMP accident model that can be used in the early systems engineering process

The first safety-related activity involves identifying a set of accidents and high level hazards. While system-level objectives define what the system *should* do, system-level hazards define what the system should *not* do (see Table 8). An accident is simply a loss that stakeholders must avoid, and a hazard is defined as "a system state or set of conditions that, together with a particular set of worst-case environmental conditions, will lead to an accident (loss)" [Leveson, 2012].

For example, accidents for a train system would include loss of human life or injury. Hazards for a simple automated train door include:

> [H-1] Doors close on a person in the doorway

> [H-2] Doors open when the train is moving or not in a station

> [H-3] Passengers or staff are unable to exit during an emergency

5.4.1 Identifying Safety Constraints

In safety-driven design, requirements typically take the form of a safety constraint. These constraints allow, control, or restrict the behavior of components as well as their interactions. For the train door hazards listed above, the associated safety constraints would be:

> [SC-1] Doors must not close when a person is in the doorway

> [SC-2] Doors must not open when the train is moving or not in a station

> [SC-3] Train must allow passengers/staff to exit during an emergency

The safety constraints are simply written in such a way that prevents the associated hazard from occurring. Much like the system-level safety constraints depend on identifying the appropriate system-level hazards, identifying lower-level safety constraints depends on identifying lower-level hazardous scenarios and causal factors.

That is, just as the system-level safety constraints are written to prevent the occurrence of a hazard, so are the lower level constraints written to prevent the occurrence of hazardous scenarios and associated causal factors. The safety-related constraints and requirements will only be complete to the extent that the analysis used to identify causal factors is complete. Developing

refined requirements and constraints, then, is made possible by rigorously defining the control structure (via methods described in Section 5.2) and then identifying potentially hazardous scenarios (via methods described in Section 5.3).

For example, a hazardous scenario might arise if the system does not satisfy Completeness Criteria for Individual Control Loops (Section 5.3.1). The train door may close on a passenger if it cannot sense the presence of passengers in the doorway. In this case, the Observability Condition is unsatisfied, and refined requirements related to the first hazard involve sensing and detecting the presence of passengers in the doorway.

Another issue might arise if the train door can be controlled by both the train conductor and automated systems. The conductor and automation might issue conflicting "Close Door" commands and violate Coordination and Consistency postulates (Section 5.3.3). Refined requirements would then constrain the behavior of the door whenever multiple agents issue simultaneous commands.

The following sections present an example of the approach used to identify scenarios, applied to an important concept being developed for the next generation of air traffic management. Section 7 then uses those results to systematically identify requirements and constraints that will prevent these scenarios from occurring.

5.4.2 Generating the Control Structure

Developing the control structure takes the following steps. The first step is developing a control model from the ConOps, using the theory and process described in Section 5.2. That is, the theory and process developed in Section 5.2 results in a model that describes the *structure* of the system and the *relationships* between components.

These types of information are consistent with a general definition of architecture given in Section 7. However, a control structure provides additional information about the relationships between components. These relationships are based on systems theory, where the control structure is represented by a hierarchical structure and each level imposes constraints on the activity of the level beneath it. Therefore, a control structure contains not only general information exchanges but also hierarchical and authority relations, control roles, and safety-related responsibilities.

For example, a general[16] architectural relationship might state that Component X exchanges data Y with Component Z. However, in a control structure there is additional information, which might state that Component X enforces constraints on Component Z via actuator signal Y. Component X is therefore at a higher level in the system hierarchy.

The theory in Section 5.2 and demonstration in the following section represent a process for developing the initial system control structure. However, this initial control structure is based on an informal ConOps that may have flaws or missing information, and thus the control structure itself may have flaws. That is, because the analysis in Section 6 focuses on identifying causal factors and missing information, the control models also contain missing or potentially

---

[16] See the definitions of a general architecture in Section 2.1.

hazardous relationships. As will be shown shortly, these causal factors can be used to further develop and/or modify the system control structure.

The resulting control model is then refined using the scenarios and causal factors identifying using the theory in Section 5.3. These scenarios and causal factors can be used in at least two ways to refine or modify the system control model.

The first way to refine the control structure is by developing requirements and additional constraints on component behavior. Just as the system safety constraints are worded to prevent the occurrence of system-level hazards, the refined safety constraints are intended to prevent the occurrence of lower level scenarios and causal factors.

The previous section described safety requirements and constraints, which is important in and of itself. The constraints also help define the system control structure. For example, if the analysis identifies a missing feedback link from a lower level component to a higher level component, then an additional requirement would add the feedback. This additional requirement (or constraint) represents an update to the control structure.

The second way to use, and hopefully prevent, the hazardous scenarios and causal factors is to change the control structure itself. Recall that a hierarchical control system consists of control agents, controlled processes, constraints going down the hierarchy in the form of control actions, and information going up the hierarchy in the form of feedback (see Figure 6). A modification of the control structure, then, consists of modifying the order of the hierarchy itself (i.e. the authority that one component has over another), the available control actions, and the available feedback.

Consider again the train door example, where the train door can be controlled by both the train conductor and automated door controller. The previous section described a potential safety constraint regarding order of priority if there is a conflict in commands between the two door control agents.

That example, and the ensuing requirements, assume some existing control structure (that both the conductor and autonomous system can open or close the door). Based on this control structure, there exist several possible requirements to deal with the case of conflicting commands. Each of these requirements has a tradeoff.

An alternative approach is to *modify the control structure itself*. In the train example, an alternative structure might disallow simultaneous door commands altogether. One control structure would eliminate manual control via the train conductor entirely and use only the automated system (or vice versa).

Clearly, modifying the control structure also has tradeoffs. It is often the case, however, that a change in the control structure is simpler and more effective than the alternative, which is attempting to derive detailed safety requirements based on an existing structure (and its associated hazardous scenarios). Section 7 presents such an example in the air traffic control domain.

5.4.3   Summary

This section has described the theoretical underpinnings of a new approach to concept development and safety-driven design. Following this theoretical development, the section proceeds with the proposed approach.

The process consists of identifying the control concepts within a ConOps, generating an initial control model, and then interrogating that model in order to identify hazardous scenarios. Once the initial control concepts and hazardous scenarios are identified, this information can then be used to generate safety-related requirements and constraints and to modify or refine the system control structure.

The following sections describe the process, applied to an example in the air traffic control domain.

6.   Application of STECA Approach

Recall the goals of this research from the previous sections. The first goal is to develop rigorous, systematic tools for the analysis of future concepts in order to identify hazardous scenarios and undocumented assumptions. The second goal is to extend these tools to assist stakeholders in the development of concepts using a safety-driven approach. Both goals especially apply to systems where the tradespace includes human operation, automation or decision support tools, and the coordination of decision making agents.

Systems-theoretic early concept analysis (STECA) has been applied to the Trajectory-Based Operations (TBO) concept being developed by the United States Federal Aviation Administration as part of the NextGen air traffic management modernization program. After a brief description of this concept, the following sections present the application to the case study, showing how the objectives can be achieved in a systematic, rigorous fashion.

6.1   *Trajectory-Based Operations*

The United States air transportation is already under stress, and demand in aircraft operations is expected to increase significantly in the next decade and beyond [FAA, 2013c]. In addition, there are growing concerns about air transportation's effect on the environment and national security. It is assumed in the aerospace community that current technologies and procedures in the national airspace cannot meet these increasing demands; therefore, the United States is creating the Next Generation Air Transportation System (NextGen) air traffic management modernization program. The goals of NextGen are to expand capacity, ensure safety, protect the environment, and grant flexibility and equity to airspace users.

The TBO Concept of Operations proves to be a useful case study for this report. There are several reasons it is compelling as a case study, in addition to the technical reasons to be described momentarily and their relationship to the literature gap identified in Section 3. First, TBO is an important real-world problem, and its successful implementation will have a significant impact on tomorrow's airspace. Past attempts at modernizing the national airspace have failed, in part due to the ineffectiveness or lack of tools necessary to develop new technologies and procedures [Cone, 2002; Office, 1986].

Second, a professional working group has conducted a preliminary hazard analysis of the TBO concept, and they used the TBO ConOps as the primary source of information for their analysis

[JPDO, 2012]. The existence of results developed using the traditional approach provides a basis on which to compare the results of this research.

Trajectory-Based Operations (TBO) is a shift from the current Air Traffic Management (ATM) and control strategy of clearance-based operations, intended primarily to increase capacity and improve efficiency. Today's operations rely on relatively little automation, in comparison to the TBO framework where aircraft will follow four dimensional paths, called trajectories, which are computed by autonomous systems and decision support tools (DST) [JPDO, 2010]. When fully realized, these trajectories will represent an aircraft's gate-to-gate movement and will be the basis for Air Traffic Control (ATC) and Air Traffic Management (ATM) that focuses on traffic flow and airspace use and autonomy of individual aircraft. The primary themes of TBO are: moving from clearance-based to trajectory-based airspace management, increasing reliance on automation and decision support tools, and distributing traffic management responsibilities throughout the system.

A key term in TBO is the four dimensional trajectory, or 4DT, which defines the aircraft in 3-dimensional space and time and is described in the list below. TBO uses the 4DT "to both strategically manage and tactically control surface and airborne operations" [JPDO, 2011]. The 4DT represents not only the aircraft's current state but also its intent, or where it will be in the future in both space and time. There are several other key terms that are listed here for reference:

    (1) 4DT - Four dimensional trajectory, defined laterally and longitudinally by latitude and longitude, vertically by altitude and with time. Surface movement is a 3DT—lateral, longitudinal, and time [JPDO, 2011].

    (2) RNP - Required Navigation Performance, describes an aircraft's ability to follow a ground track and/or vertical profile to within some specified tolerance in nautical miles or feet (± TBD NM, ± TBD ')

    (3) RTP - Required Time Performance, describes an aircraft's ability to reach a way-point within some specified window of time

    (4) Conformance - Monitoring of the aircraft's position, altitude, and time performance against the agreed-upon 4DT. Monitoring is against performance requirements for the flight maneuver or surface movement. Conformance monitoring occurs both in the air and within ground automation. Alerts are generated if the aircraft is not meeting its 4DT performance [JPDO, 2011].

These definitions allude to the fact that TBO relies not only on accurate state information but also on accurate prediction and navigation capability. In addition to these important considerations, note again that one of the goals of TBO[17] is to grant increased autonomy, flexibility, and equity to both individual users and operation centers. All of these factors represent a major shift in (1) the types of information gathered and exchanged, (2) technologies used to generate and transmit this information, (3) the roles and authority of the various actors.

In addition to the major technological and administrative challenges facing the development and implementation of NextGen, there are general research objectives that this report is intended to address. There is often informality and lack of rigor often present during concept generation, and

---

[17] This goal of TBO is also an important goal of NextGen in general.

Section 5.1 describes some of the challenges specific to generating and analyzing a Concept of Operations. The TBO Concept of Operations [JPDO, 2011] shares many of the problems typically found in a ConOps: lack of a specification or requirements, prevalence of undocumented or implicit assumptions, and description of the concept using informal text or graphics.

6.2 Analysis of TBO

The basic steps of the proposed approach described in Section 5 are shown again in Figure 13. This section demonstrates the two highlighted boxes in Figure 13, namely, the identification of hazards and safety constraints. This section also develops the general air traffic management control hierarchy, using concepts and theory from the previous section.
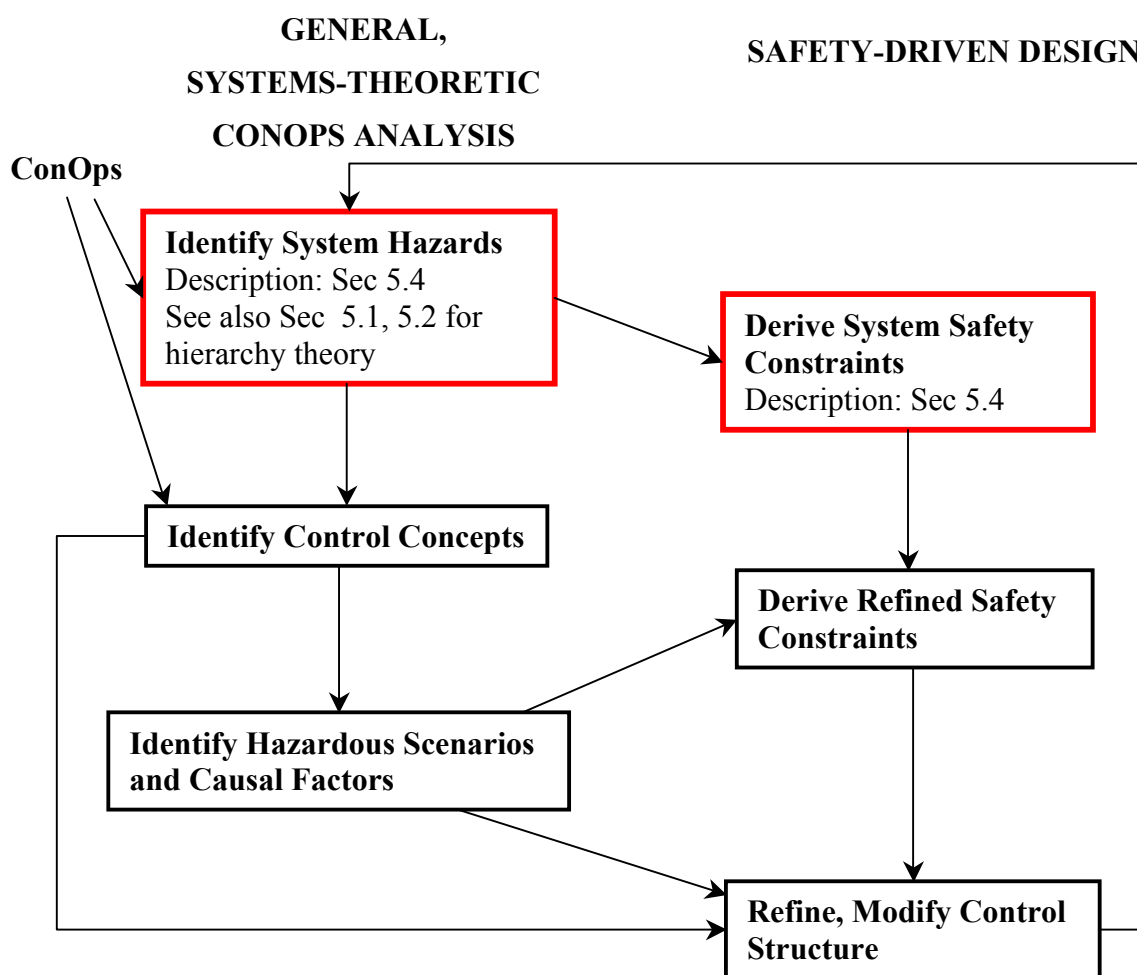
Figure 13. Methodology - Top-Level Systems Engineering

Table 7 elsewhere depicts the relationships between safety-driven design activities and their counterparts in general systems engineering. The first safety-related activity involves identifying a set of accidents and high level hazards (see Table 8 for definitions).

For example, in air traffic management, the accidents would be aircraft loss, equipment damage, and injury or loss of life. A few example hazards that could lead to these accidents are listed below. This list represents a sub-set of hazards used for the certification of TCAS [Leveson and Reese, 1994] and the analysis of several aeronautical applications [Fleming et al., 2013].

(1) [H-1] Aircraft violate minimum separation (LOS or loss of separation, NMAC or Near midair collision)
(2) [H-2] Aircraft enters uncontrolled state
(3) [H-3] Aircraft performs controlled maneuver into ground (CFIT, controlled flight into terrain)

Using these hazards, then define a set of safety-related requirements (constraints) and responsibilities. The system-level safety constraints relate directly to the hazards; conforming to the safety constraints should prevent the associated hazard from occurring. The following constraints are based on the three hazards above.

(1) [SC-1] Aircraft must remain at least TBD nautical miles apart en route[18] ↑ [H-1]
(2) [SC-2] Aircraft position, velocity must remain within airframe manufacturer defined flight envelope ↑[H-2]
(3) [SC-3] Aircraft must maintain positive clearance with all terrain (This constraint does not include runways and taxiways) ↑[H-2]

The next step involves defining a high-level system control structure. As described in Section 5, in STAMP the system is represented by a hierarchical, functional control structure. In a Level of Decision Complexity hierarchy, higher level responsibilities are concerned with slower aspects of the system operation, the disturbances take place at a lower frequency, and the dynamics of concern are slower.

Among the hazards and constraints listed above, the lowest level function involves aircraft attitude control ([H-2],[SC-2]). That is, control of the aircraft's speed, heading, and altitude occurs on the (relatively) shortest time scale, involves the least amount of uncertainty, and requires the least amount of information. This level of control occurs via changes in thrust and manipulation of aircraft control surfaces. The next level involves avoiding terrain ([H-3],[SC-3])—the aircraft must use its control systems, but there is added complexity associated with identifying both the terrain and potential changes in trajectory in order to avoid those obstacles. Not only are the control decisions more complex, but also this function occurs on a slower time scale than the manipulation of aircraft thrust and control surfaces. Finally, the top level of the hierarchy involves the separation of two or more aircraft ([H-1],[SC-1]). Separation of aircraft represents the highest level of complexity among the listed hazards. Separation assurance necessitates not only knowledge and prediction of multiple aircraft states but also the identification of, and selection among, multiple actions required to avoid a conflict.

The basic functions required to safely manage the airspace and prevent the above hazards from occurring, then, are route planning[19] and piloting[20]. The route planning function must provide

---

[18] Similar constraints could be developed for other phases of flight.

conflict-free trajectories and sequence the flow of aircraft, while the piloting function must navigate the aircraft along its assigned path. From a hierarchical, control-theoretic perspective, this example system takes the structure shown in Figure 14. The safety-related responsibilities are also stated in that figure. A hierarchy of route planning, guidance, and control is a typical decomposition in engineered systems [see for example Leonard et al. 2007].

These functions and responsibilities are intentionally solution neutral and are applicable to any air traffic management concept. That is, these functions are not (yet) assigned to a human or computer, aircraft or air traffic controller, or any combination of these and other potential solutions. While the TBO ConOps has made some assumptions about these assignments, the level of generality, neutrality, and abstraction in Figure 14 is important in systems engineering, particularly during the architecting phases.

---

[19] "Route Planning" is intended to be a very high level, abstract term. This is the general term for the entities responsible for managing air tra?c, making ultimate decisions on aircraft routes, identifying andsolvingconflicts, etc. TraditionallythishasbeentheresponsibilityofATCandassociateddecision support tools; in the future this responsibility may be delegated to operators, flight crews, and/or autonomous systems.

[20] "Piloting" is also intended to be a very high level, abstract term. This is the general term for the entities responsible for navigating and controlling aircraft. Traditionally this has been the responsibility of the flight crew, in conjunction with flight management systems. In the future, the 'pilot' could be any combination of flight crew and airborne automation, or even ground-based crews or automation.

Function	Safety-Related Responsibilities

Route Planning
- Provide conflict-free clearances & trajectories
- Sequence the flow of aircraft
- Trajectory management

Piloting
- Navigate the aircraft
- Provide aircraft state information to route planning func
- Avoid conflicts with other aircraft, terrain, weather
- Ensure that trajectory is within aircraft flight envelope

Aircraft
- Provide lift
- Provide propulsion (thrust)
- Orient and maintain control surfaces

Environment

Figure 14. High Level Control Structure & Responsibilities

These activities provide a reference and context upon which one can develop a model of the system using control- and systems-theoretic processes. In safety-driven design, the development of the model and subsequent analysis of the model should ultimately be traceable to system hazards and high level safety-related responsibilities and functions. Again, these activities and artifacts closely parallel general systems engineering.

In fact, identifying hazards and safety-related responsibilities should be done in conjunction with identifying system goals and requirements. This section, along with Section 5, describes the safety-driven identification of a system control structure that is based on control and systems theory. There is nothing preventing the use of a similar approach for generating a control structure that satisfies other system properties, in addition to safety.

At this early stage, the hierarchy should be (a) sufficiently general, in order to avoid overly constraining potential design solutions and (b) account for the system-level losses that stakeholders want to avoid, hazards associated with those losses, and constraints that will prevent those hazards. That is, the hierarchy should be relatively simple and concise, and yet these steps are important because they provide the framework and context for developing and analyzing the model. The following sections describe model generation (6.3 and analysis (6.3.1 for TBO Conformance Monitoring, one important aspect of the overall TBO concept.

6.3   Model Generation

Recall from Section 5, (see also Figure 15) the description and theory used to generate a model from textual or graphical descriptions of a concept. That theory is now applied to one section of

the TBO ConOps, which describes the Conformance Monitoring function, while Appendix A contains further example analyses of TBO.
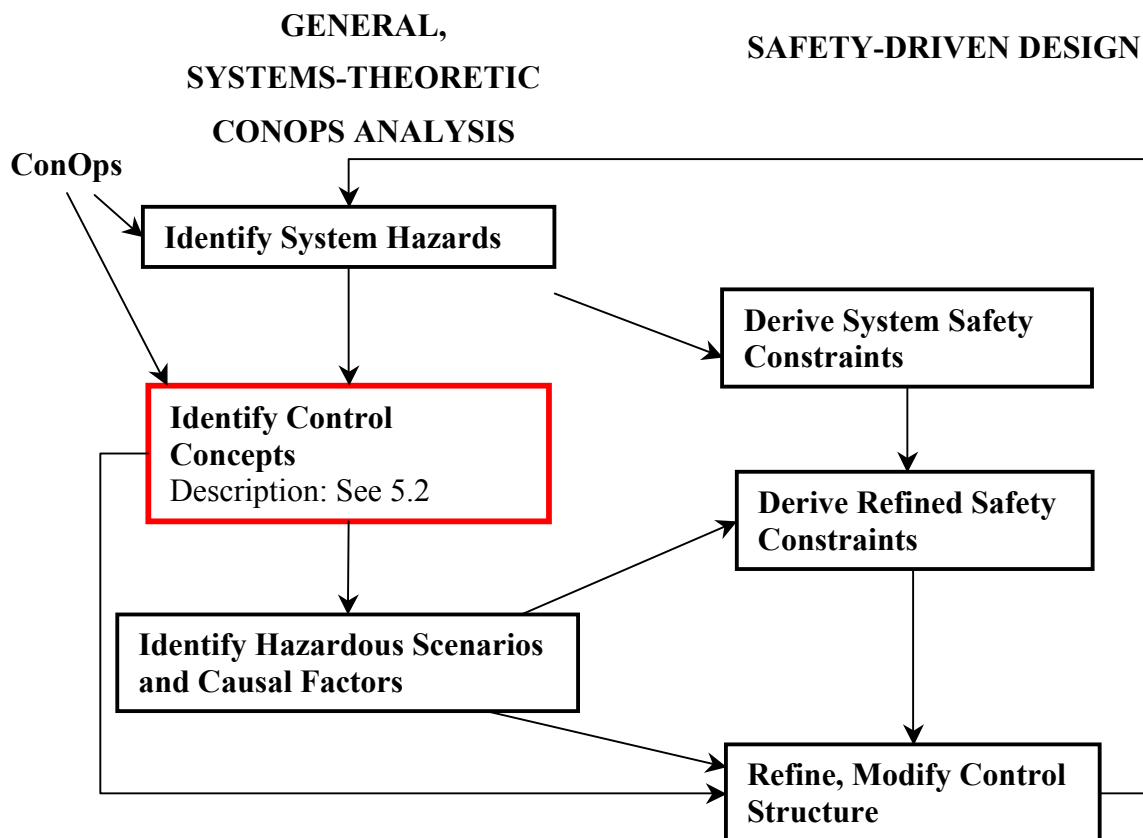
Figure 15. Methodology - Identifying Control Concepts

In the TBO ConOps [JPDO, 2011], there is a section dedicated to conformance monitoring, which is the degree to which an aircraft follows its agreed-upon trajectory. This example is intended to show how these control-theoretic concepts can be used to (1) query a certain aspect of a concept, it could be a sentence, paragraph, or figure and then (2) to use the resulting information to build a control-theoretic model of the system. The following quote is one of the first sentences dedicated to conformance monitoring:

> ∃ -1   *TBO conformance is monitored both in the aircraft and on the ground against the agreed-upon [trajectory]. In the air, this monitoring (and alerting) includes lateral deviations...(actual lateral position compared to intended position), longitudinal based on flight progress in the FMS [aircraft software], vertical based on altimetry, and time from the FMS [aircraft software] or other "time to go" aids. [JPDO, 2011]*

To begin, the analyst must ask: What is the primary source, subject, or actor in the text, and in what way does this source relate to control theory? The quoted text describes conformance, or *conformance* monitoring.

Next, what is the source's role in control theory? Conformance monitoring acts as a *sensor*, and in this text there appear to be two versions of the sensor: one in the aircraft and another on the ground. Of the three generic roles that a sensor can take in the proposed framework, the conformance monitoring sensor provides two. See Table 9 for a summary of this brief example analysis, and note that a separate model will be developed for the ground conformance monitor shortly.

With this high-level picture of conformance monitoring and its role in the context of control theory, the rest of the information from the quoted text can be used to fill out further details about the control model. From the text we can identify that conformance monitoring has directly to do with identifying an aircraft's position (latitude, longitude, altitude, at a given time). The text also suggests that this position will be compared with the aircraft's intended position. A sensor is used to inform a controller's process model. Therefore, it can be assumed that these position variables are related to a process model.

Table 9. Example Analysis of Text—TBO Conformance Monitoring

Subject	Conformance monitoring, Air automation
Role	Sensor
Behavior Type	Transmits binary or discretized state data to controller (i.e. measures behavior of process relative to thresholds; has algorithm built-in but no control authority)
	Synthesizes and integrates measurement data
Context	This is a decision support tool that contains algorithms to synthesize information and provide alerting based on some criteria.

Next, the conformance monitor uses the FMS (aircraft software) and associated equipment to measure these process model variables. Finally, from the high level control structure and responsibilities, shown in Figure 14, it can be inferred from the text that the airborne conformance monitor is related to the control function of piloting the aircraft.

Table 10 shows these results in tabular form, relating to each control loop element necessary for ensuring control and coordination. Notice that the table is incomplete. Although much of the missing information can be inferred, particularly by an expert on this system, the table cannot and should not be completed without further investigation into the rest of the ConOps or other NextGen documentation. Also, it is important to bear in mind that this incompleteness does not imply that the ConOps is "incomplete". Rather, this is simply the first step in the systematic, recursive process and completeness can only be considered when this process has been applied over the entire concept, or system boundary.

Table 10. Preliminary Control Model of Conformance Monitor Example

See Figure 9		**Description**
1.	Controller	Piloting function
2.	Actuator	
3.	Cntl'd Process	Aircraft

4.	Sensor	Altimeter, FMS, Aircraft conformance monitor
5.	Process Model	Intended latitude, longitude, altitude, time; Actual latitude, longitude, altitude, time
6.	Cntl Algorithm	
7.	Control Actions	
8.	Controller Status	
9.	Control Input	
10.	Controller Output	
11.	External Input	
12.	Alt Controller	
13.	Process Input	
14.	Proc Disturbance	
15.	Process Output	

Figure 16 graphically depicts how the natural language text can be mapped onto a control loop. While a database format of Table 8 and Table 9 assists in data storage, analysis, and retrieval, a graphical representation improves a user's ability to visualize the control-theoretic elements of the concept. Both formats, databases and graphical control loops, lend themselves to a formalism that allows for automatic checking of consistency and completeness.

Finally, quote ⌐ -1 also mentions conformance monitoring on the ground. Though its role is notionally the same, ground monitoring represents a different source of information. The model should reflect this difference. Because the rest of the quoted text describes conformance monitoring in the air, little can be inferred about the ground version of conformance monitoring and thus the analyst must look for additional detail elsewhere in the ConOps. The initial model for ground-based conformance monitoring has very little detail and is shown in Table 11.

Table 11. Initial Control Model of Ground Conformance Model

Subject	Conformance monitoring, Ground automation
Role	
Behavior Type	
Context	

⇓

1.	Controller	
2.	Actuator	
3.	Cntl'd Process	Airspace
4.	Sensor	

5.	Process Model	
6.	Cntl Algorithm	
7.	Control Actions	
8.	Controller Status	
9.	Control Input	
10.	Controller Output	
11.	External Input	
12.	Alt Controller	
13.	Process Input	
14.	Proc Disturbance	
15.	Process Output	

Consider the next sentence in the TBO ConOps:

> ❥-2 *Independent of the aircraft, the ANSP uses ADS-B position reporting for lateral and longitudinal progress, altitude reporting for vertical, and tools that measure the time progression for the flight track. Data link provides aircraft intent information. Combined, this position and timing information is then compared to a performance requirement for the airspace and the operation. ...precision needed...will vary based on the density of traffic and the nature of the operations. [JPDO, 2011]*

Prior to the above text, the analysis thus far only reveals that there is a ground-based conformance monitor, it performs the role of sensor, and the controlled process is the aircraft. The quote reveals several additional things about ground-based conformance monitoring: the behavior type, the type of sensing used by the ANSP, the process model states to be used, as well as additional sources of information used by the ground (ANSP) entity. See Table 4.4 for a summary of this brief example analysis, and note that a separate model is developed for the ground conformance monitor.

Table 12. ANSP/Ground—TBO Conformance Monitoring

Subject	Ground automation
Role	Sensor
Behavior Type	Synthesizes and integrates measurement data
Context	This is a decision support tool that contains algorithms to synthesize information and provide alerting based on some criteria.

The information in Table 12 is similar to that in Table 9. However, the information contained in this particular quoted text is much different than that in ❥-1, and the resultant control model is much different. Table 13 shows that the ground-based conformance monitor may be using a different set of tools to monitor conformance than airborne conformance monitoring, and the ground-based monitor also has a different process model and controlled process.

Table 13. Preliminary Control Model of Ground Conformance Monitor

See Figure 9		Description
1.	Controller	Piloting function
2.	Actuator	
3.	Cntl'd Process	Aircraft
4.	Sensor	Altimeter, FMS, Aircraft conformance monitor
5.	Process Model	Intended latitude, longitude, altitude, time; Actual latitude, longitude, altitude, time
6.	Cntl Algorithm	
7.	Control Actions	
8.	Controller Status	
9.	Control Input	
10.	Controller Output	
11.	External Input	
12.	Alt Controller	
13.	Process Input	
14.	Proc Disturbance	
15.	Process Output	

Adapted from [JPDO, 2011]:
conformance is monitored both in the *(1.),(3.) aircraft* and on the ground against the agreed-upon [trajectory]. In the air, this monitoring (and alerting) includes lateral deviations...*(5.) actual lateral* position compared to *(5.) intended position*, longitudinal based on flight progress in the *(4.) FMS [aircraft software]*, *vertical based on altimetry*, and time from the FMS [aircraft software] or other "time to go" aids.

(1.,5.)   *(4.)*

1. Controller
- Piloting Function

5. Process
Model
(x,y,h,t...)

(3.)

4. Sensor -
Altimeter, FMS, aircraft conformance monitor

3. Controlled Process
-Aircraft

Figure 16. Graphical Control Model of Airborne Conformance Monitor

72

Adapted from [JPDO, 2011]:
Independent of the aircraft, the *(1.) ANSP* uses *(4.) ADS-B* position reporting for
(5.) lateral and longitudinal progress, *(5.) altitude* reporting for *(5.) vertical*,
and *(4.) tools* that measure the *(4.) time* progression for the flight track. *(8.)*
Data link provides aircraft *(5.) intent* information. Combined, this position and
timing information is then compared to a *(5.) performance requirement* for the
airspace and the operation. ...precision needed...will vary based on the *(5.)*
density of traffic and the *(5.) nature of the operations*.

8. Datalink *(1.,5.,8.)*   *(4.)*

1. Controller
- ANSP/Ground

5. Process Model (x,y,h,...)

(3.)

4. Sensor -
ADS-B, Alt Rep,
time, grd
conformance
monitor

**3. Controlled
Process**
-Piloting Function
& Aircraft
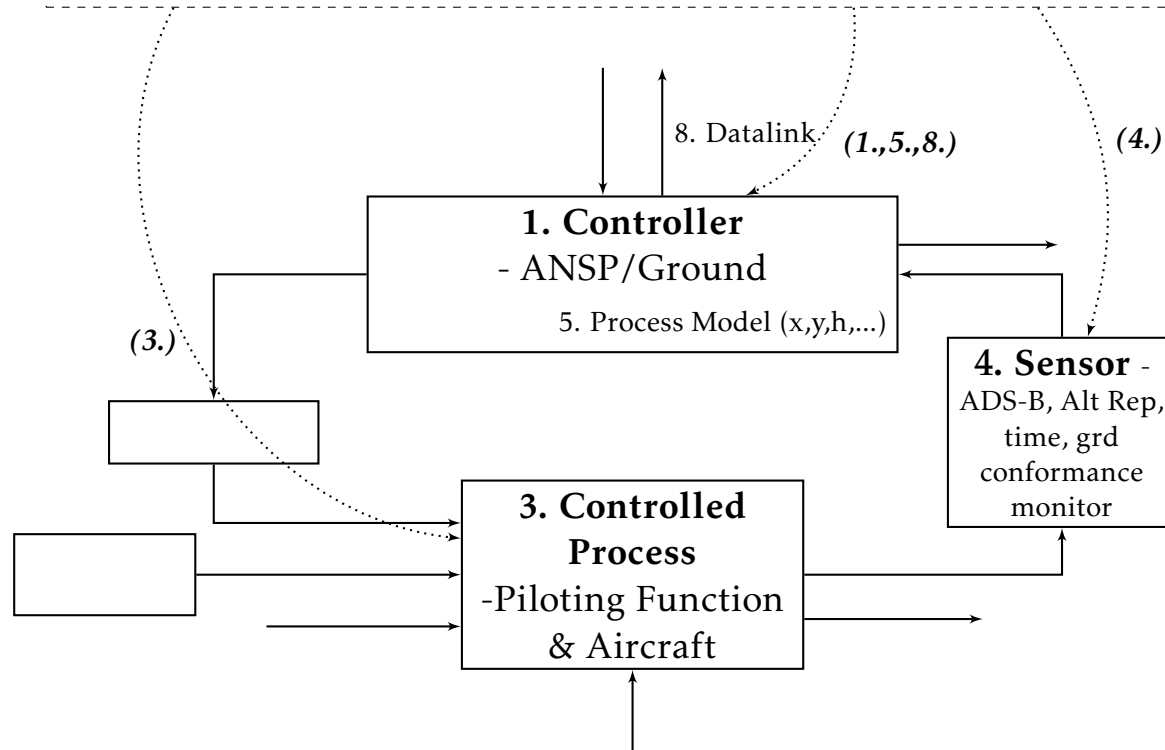
Figure 17. Graphical Control Model of Ground Conformance Monitor

73

6.3.1 Recursive Application of Method - Conformance Monitoring

The following presents a recursive application of the method over the remainder of the section in the TBO ConOps dedicated to conformance monitoring. After each iteration the control model is updated, and the updated aspects of the model are depicted as **bold and underlined**, while the parts of the model from previous iterations are depicted using standard font.

The following quote adds more detail about performance requirements and their relation to conformance monitoring. The updated control model is shown for the ground automation only in Table 14, although similar results hold for the airborne automation.

> *ꟻ -3 In framing the required performance...,TBO recognizes that traffic density drives needed performance. There may be departures where a lateral precision of RNP 0.3 is required close in to the airport, and where time is measured in seconds. RTP is used as a tool to separate crossing traffic, and where vertical altitude restrictions are necessary. All of these factors must be considered in defining the parameters for conformance monitoring. [JPDO, 2011]*

Table 14. Updated Control Model for ꟻ -3

See → ꟻ -3 of C$_{TBO}$	
Subject	Conformance monitoring, Ground automation
Role	Sensor
Behavior Type	Synthesizes and integrates measurement data
Context	This is a decision support tool that contains algorithms to synthesize information and provide alerting based on some criteria.

⇓

1.	Controller	Ground ≡ ANSP
2.	Actuator	
3.	Cntl'd Process	Airspace
4.	Sensor	ADS-B, altitude reporting, "tools", Data-link
5.	Process Model	All (all aircraft under ANSP jurisdiction) Intended lat, long, alt, time; All Actual lat, long, alt, time; traffic density; operation type; performance requirement **RNP, RTP**
6.	Cntl Algorithm	**Crossing vs flow traffic**
7.	Control Actions	
8.	Controller Status	
9.	Control Input	
10.	Controller Output	

11.	External Input	
12.	Alt Controller	
13.	Process Input	
14.	Proc Disturbance	
15.	Process Output	

The following information (ℸ -3) from the TBO ConOps describes the predicted aircraft state information in slightly different terms than the previous descriptions. Therefore Table 15 includes the new text along with the text it replaced.

ℸ -4 Conformance monitoring has an expected ground track, climb performance (based on known aircraft type, weight, and preferred profile), and time performance. In conformance monitoring, the aircraft is on a closed trajectory. [JPDO, 2011]

Table 15. Updated Control Model for ℸ -4

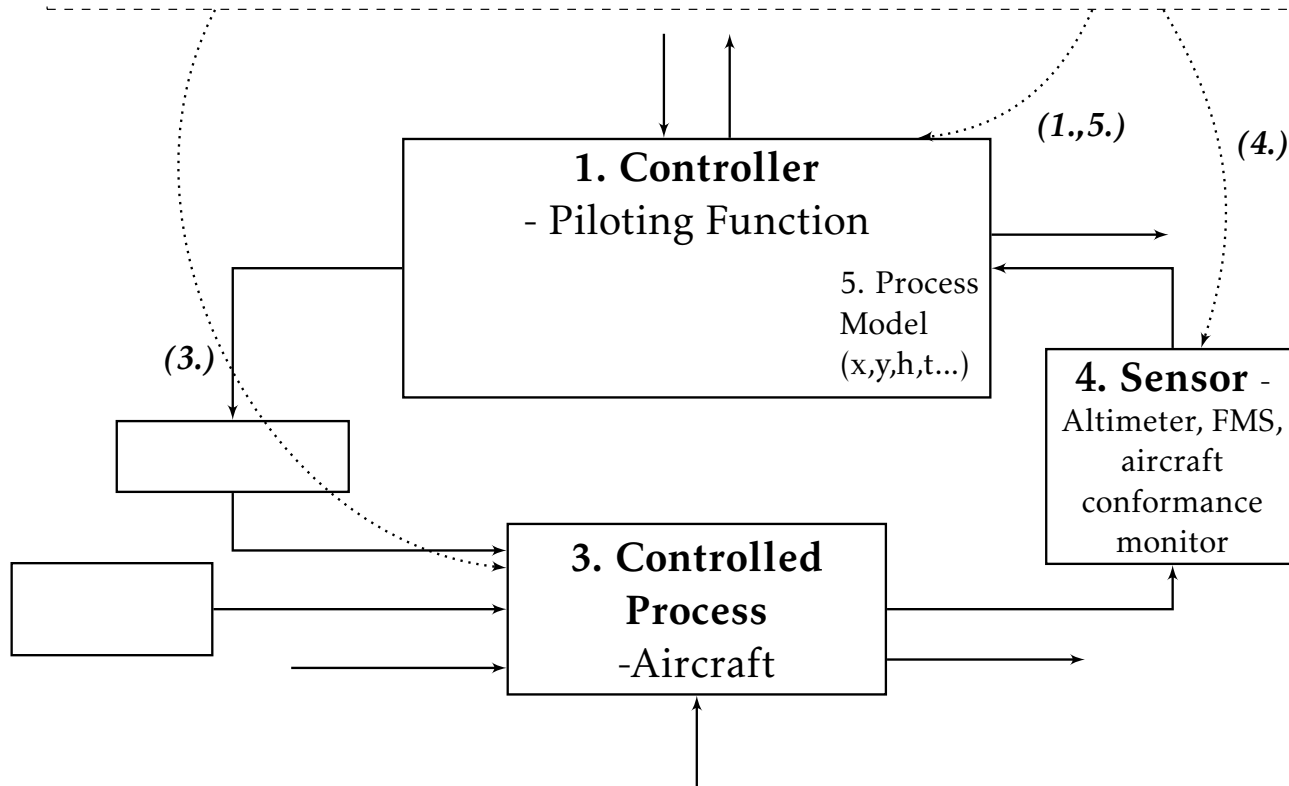| See → ℸ -4 of C_TBO	
Subject	Conformance monitoring, Ground automation
Role	Sensor
Behavior Type	Synthesizes and integrates measurement data
Context	This is a decision support tool that contains algorithms to synthesize information and provide alerting based on some criteria.

⇓

1.	Controller	Ground ≡ ANSP
2.	Actuator	
3.	Cntl'd Process	Airspace
4.	Sensor	ADS-B, altitude reporting, "tools", Data-link
5.	Process Model	All (all aircraft under ANSP jurisdiction) Intended **ground track, climb performance, time performance** ~~lat, long, time~~; All Actual **ground track, climb performance, time performance** ~~lat, long, time~~; traffic density; operation type; performance requirement RNP, RTP; aircraft type, weight, performance **profile** ~~requirement~~
6.	Cntl Algorithm	Crossing vs flow traffic
7.	Control Actions	
8.	Controller Status	

9.	Control Input	
10.	Controller Output	
11.	External Input	
12.	Alt Controller	
13.	Process Input	
14.	Proc Disturbance	
15.	Process Output	

The next quote from the TBO ConOps (‫ ‬-5) describes separation in terms of the overlap between the protected airspace of two or more aircraft. The protected airspace is described as an elliptical shape, and the size of these ellipses is dependent on the certainty of surveillance and navigation information. Importantly for TBO, this characterization of protected airspace is much different than the current paradigm of protected airspace, where every aircraft is laterally separated by a fixed distance or vertical separation, depending on the operation and type of surveillance [Ray, 2014].

In the traditional approach, protected airspace is a disc with a radius equal to the lateral separation and height equal to the vertical separation. Table 16 includes updated information about potentially unsafe actions for both the general functions in air traffic management as well as updated information for the process model and control algorithm. Table 16 describes only the ground automation portion of the system, but similar updates are included for airborne automation.

> ‫ ‬-5   As the aircraft approaches level-off and cruise, the shape of the protected
> airspace morphs into more of an elliptical 3-D shape, where the aircraft is
> positioned in the narrow end of the elliptical shape, with the wake vortex "tail"
> as its aft bound and vertical, lateral, and longitudinal uncertainty defining the
> flexible airspace. No two elliptical shapes can overlap if separation is to be
> assured. In this case, Aircraft A and Aircraft B have crossing trajectories.
> Aircraft A's protected space is smaller because it has less uncertainty than
> Aircraft B. The trailing area of protection may reflect wake turbulence
> requirements. The lateral protection is the uncertainty in navigation performance,
> while the leading distance along the flight path represents the time uncertainty. In
> level flight, the vertical altitude dimension is quite small. [JPDO, 2011]

Table 16. Updated Control Model for ‫ ‬-5

See → ‫ ‬-5 of c_{TBO}	
Subject	Conformance monitoring, Ground automation
Role	Sensor
Behavior Type	Synthesizes and integrates measurement data

Context	This is a decision support tool that contains algorithms to synthesize information and provide alerting based on some criteria. **Potential unsafe control action for trajectory generation function: Approving a 4DT that will lead to LOS or not modifying a 4DT that will overlap Potential unsafe control action for piloting function: Aircraft is following a 4DT that will lead to LOS**	

⇓

1.	Controller	Ground ≡ ANSP
2.	Actuator	
3.	Cntl'd Process	Airspace
4.	Sensor	ADS-B, altitude reporting, "tools", Data-link
5.	Process Model	All (all aircraft under ANSP jurisdiction) Intended ground track, climb performance, time performance; All Actual ground track, climb performance, time performance; traffic density; operation type; performance requirement RNP, RTP; aircraft type, weight, performance profile; **Wake turbulence; Ellipse, uncertainty** (shape of conformance tolerance)
6.	Cntl Algorithm	Crossing vs flow traffic; **Overlap of 2 or more 4DTs (Ellipses)**
7.	Control Actions	
8.	Controller Status	
9.	Control Input	
10.	Controller Output	
11.	External Input	
12.	Alt Controller	
13.	Process Input	
14.	Proc Disturbance	
15.	Process Output	

The TBO ConOps also describes an alerting function for the conformance monitoring automation for both the ground controller (Ⅎ -6) and flight crew (Ⅎ -7). The basic idea of the alerting function is to allow the user - either ATC or flight crews, independently - to set alerts for measuring an aircraft's progress against its assigned 4DT. Table 17 contains the updated model information for the ground controller, developed from quote Ⅎ -6.

ꞁ -6   Alerting is triggered by automation and alerts the controller to transgression from the conformance airspace, and may be set as alerts for measuring progress. By setting progress alerts, the controller has an aid to measure progress in meeting the 4DT. [JPDO, 2011]

Table 17. Updated Control Model for ꞁ -6

See → ꞁ -6 of ℭ_{TBO}	
Subject	ATC
Role	Controller
Behavior Type	Processes inputs from sensors to form and update process model
Context	ATC measures progress via conformance monitoring automation Potential unsafe control action for trajectory generation function: Approving a 4DT that will lead to LOS or not modifying a 4DT that will overlap

⇓

1.	Controller	**ATC (controller)**
2.	Actuator	
3.	Cntl'd Process	Piloting function and aircraft, Airspace
4.	Sensor	ADS-B, altitude reporting, "tools", Data-link, **alerting automation**
5.	Process Model	All (all aircraft under ANSP jurisdiction) Intended ground track, climb performance, time performance; All Actual ground track, climb performance, time performance; traffic density; operation type; performance requirement RNP, RTP; aircraft type, weight, performance profile; Wake turbulence; Ellipse, uncertainty
6.	Cntl Algorithm	**Generate or compute conformance monitoring airspace volume [automation]; compare actual position with monitoring volume [automation]**
7.	Control Actions	
8.	Controller Status	
9.	Control Input	
10.	Controller Output	

11.	External Input	
12.	Alt Controller	
13.	Process Input	
14.	Proc Disturbance	
15.	Process Output	

Table 18 contains the updated model information for the ground controller, re-framing quote ⌐ -7 in control-theoretic terms.

> *⌐ -7   From the cockpit, the pilot can monitor performance, as well. Most of the tools are already used. Altitude alerts exist. RNP can be monitored, and the progress can be provided by the FMS. What is needed is the cockpit display of traffic information (CDTI) with tools for merging, spacing, and separation. These tools will help the pilot monitor other traffic as well as progress in meeting the 4DT. The pilot sets the alerting parameters in the respective automation.* [JPDO, 2011]

Table 18. Updated Control Model for ⌐ -7

See → ⌐ -7 of C$_{TBO}$	
Subject	Flight Crew
Role	Controller
Behavior Type	Creates, generates, or modifies control actions based on algorithm or procedure and perceived model of system Processes inputs from sensors to form and update process model Processes inputs from external sources to form and update process model
Context	Crew measures progress via conformance monitoring automation crew makes decisions for merging, spacing, separation Potential unsafe control action for piloting function: Aircraft is following a 4DT that will lead to LOS

⇓

1.	Controller	**Flight Crew**
2.	Actuator	Not specified; assume FMS and manual control
3.	Cntl'd Process	Implied; Aircraft control surfaces, thrust
4.	Sensor	**CDTI, FMS, Altitude alerts**, conformance monitoring automation

5.	Process Model	Ownship Intended ground track, climb performance, time performance; Ownship Actual ground track, climb performance, time performance; traffic density; operation type {climb,cruise,arrival}; performance requirement RNP, RTP; aircraft type, weight, performance profile; flight object
6.	Cntl Algorithm	**Implied: Other aircraft (relevant, all?)** Intended ground track, climb performance, time performance; Ownship Actual ground track, climb performance, time performance; traffic density; operation type {climb,cruise,arrival}; performance requirement RNP, RTP; aircraft type, weight, performance profile; flight object
7.	Control Actions	
8.	Controller Status	
9.	Control Input	**Not specified; related to merging, spacing, sequencing, and assuring conformance**
10.	Controller Output	
11.	External Input	
12.	Alt Controller	**ATC (ANSP)**
13.	Process Input	
14.	Proc Disturbance	
15.	Process Output	

The final text (¶-8) regarding conformance monitoring in the TBO ConOps is concerned with open versus closed trajectories. This is the essence of what conformance monitoring is intended to achieve: the determination of whether or not an aircraft is following its assigned 4DT. That is, conformance monitoring determines whether an aircraft is on an open or closed trajectory.

The language used in ¶-8 is slightly different than the previous conformance-related text (¶-1 through ¶-7). First, the text simply says "Automation", and it is referring to the more general TBO-related automation beyond the conformance-related automation. Next, the text mentions downstream conflicts and not just immediate overlaps of different trajectories; earlier quotes from the text mention only the overlap of elliptical airspace (¶-5 above). Finally, the text states that open trajectories can "even lead to a conflict requiring intervention".

Though the model developed in Table 19 does not include such explicit terms, the emphasis on "intervention" for only open trajectories implies that (1) TBO automation will generate conflict-free, closed trajectories; (2) the automation cannot handle open trajectories (stated in weaker terms in ¶-8); and (3) that intervention should be rare. Furthermore, from elsewhere in the text it is inferred that this intervention will come from ground controllers, but the possibility exists that on-board collision avoidance systems will also provide this intervention function.

Ɔ -8   It is difficult for automation to deal with open trajectories. The uncertainties that open trajectories introduce affect more than just the aircraft in question and may impact downstream flows, and even lead to a conflict requiring intervention to assure safety. [JPDO, 2011]

Table 19. Updated Control Model for Ɔ -8

See → Ɔ -8 of C$_{TBO}$	
Subject	Automation
Role	Controller
Behavior Type	Creates, generates, or modifies control actions based on algorithm or procedure and perceived model of system Processes inputs from external sources to form and update process model
Context	Ɔ -8 does not specify whether there is a human user of this automation, or the user and location (e.g. ground or flight deck). It does imply, however, that it is the general "Trajectory Generation Function" that is in control It is assumed that controller / ANSP "intervention" is intended to be rare

⇓

1.	Controller	**ANSP or Flight Crew**
2.	Actuator	
3.	Cntl'd Process	Piloting function and aircraft, Airspace
4.	Sensor	ADS-B, altitude reporting, "tools", Datalink, alerting automation, **TBO Trajectory automation**
5.	Process Model	All (all aircraft under ANSP jurisdiction) Intended ground track, climb performance, time performance; All Actual ground track, climb performance, time performance; traffic density; operation type; performance requirement RNP, RTP; aircraft type, weight, performance profile; Wake turbulence; Ellipse, uncertainty; **Downstream flows (flow model, in addition to "All Intended trajectory" model)**

6.	Cntl Algorithm	Generate or compute conformance monitoring airspace volume [automation]; compare actual position with monitoring volume [automation]; **Decision about aircraft conformance vs non-conformance [automation]**
7.	Control Actions	
8.	Controller Status	
9.	Control Input	
10.	Controller Output	
11.	External Input	
12.	Alt Controller	
13.	Process Input	
14.	Proc Disturbance	
15.	Process Output	

It is important to again emphasize that the model development process should be limited, to the extent possible, to the individual information element that is currently being analyzed. Because the proposed method is easily repeatable and should be applied to the entire set of information elements, $\jmath \in \mathbf{C}$, the control model represented by Table 19 contains only the assumptions that can be gleaned from $\jmath$ -8, while also documenting additional assumptions or inferences in the "Context" row.

(a)   (b)
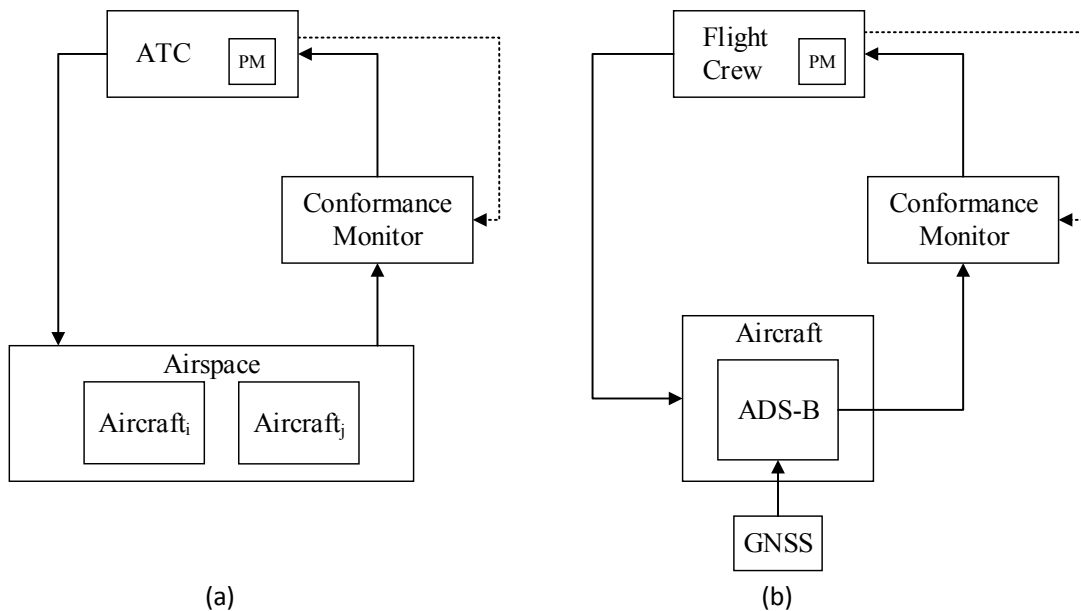
Figure 18. Individual Control Loops derived via Analysis

At this point there are two separate control models under development—one of the aircraft as a control entity and another of the ground as a control entity. Again, any distinction among pilots or aircraft avionics (e.g. FMS) as having control of the "airborne" function is unnecessary at this point in the analysis. However, the TBO ConOps development of conformance monitoring implies that pilots and air traffic controllers will still have control authority in mid- and long-term TBO systems. Figure 18 provides a high-level depiction of these separate feedback loops.

Because this research is concerned with identifying behaviors that emerge due to the interactions between and among various components, it is important to determine if and how these individual control loops relate to each other. Some of this work has already been done, as Section 6.2 laid the groundwork via a general, hierarchical decomposition of the functions required to maintain safe airspace.

These distinct control models, then, can be synthesized via a mapping to the functional control structure and hierarchical safety responsibilities derived earlier. In particular, the TBO ConOps states elsewhere that the "ANSP's authority over the airspace and the flight crew's authority over the aircraft's trajectory do not change" [JPDO, 2011]. It is therefore appropriate and intuitive to map the ground-based conformance monitoring loop to the `Trajectory Management Function' and the airborne loop to the `Piloting Function' of Figure 14.

Figure 19 depicts the control structure that results from analyzing only the section dedicated to conformance monitoring in the TBO ConOps, with the rest of details identified in Table 10– Table 19. Appendix A, contain the rest of the model generation results for other aspects of the TBO Concept of Operations.

Following is a description of each element in the synthesized model (Figure 19). GNSS, ADS-B, CDTI, and FMS are not included in the development here because they have already been developed thoroughly elsewhere.

Though conformance monitoring functions clearly have a model of the process and an algorithm for determining if an aircraft conforms to its prescribed trajectory, conformance monitoring is not a controller. Rather, conformance monitoring is a lower level function that maps refined state information into more abstract state information. That is, conformance monitoring takes detailed surveillance information (among other things), and in its simplest form outputs a binary datum to a higher level control function. The control functions, which could be other forms of automation or human operators, ultimately have to decide whether an aircraft conforms and then determine an action based on this understanding of aircraft conformance. It is for these reasons, as well as the preceding analyses (Table 10-Table 19), that the conformance monitors are depicted as sensors in Figure 19.

Because the conformance monitors are modeled as sensors, the control functions in Figure 4.7 may be modeled with relatively high level, abstract process models and algorithms.

Route, Trajectory Management Function

GROUND (ANSP / ATC)

CA$_G$ PM$_G$

Alert parameter (G)

Data Link

Conformance Monitor [Gnd]

Altitude Report

{4DT} (Intent)

Piloting Function

AIR (Flight Crew)

CA$_A$ PM$_A$

FMS; Manual

Alert parameter (A)

Conformance Monitor [Air]

CDTI

Aircraft

ADS-B

{x,y,h,t}

{x,y,h,t}

{h}
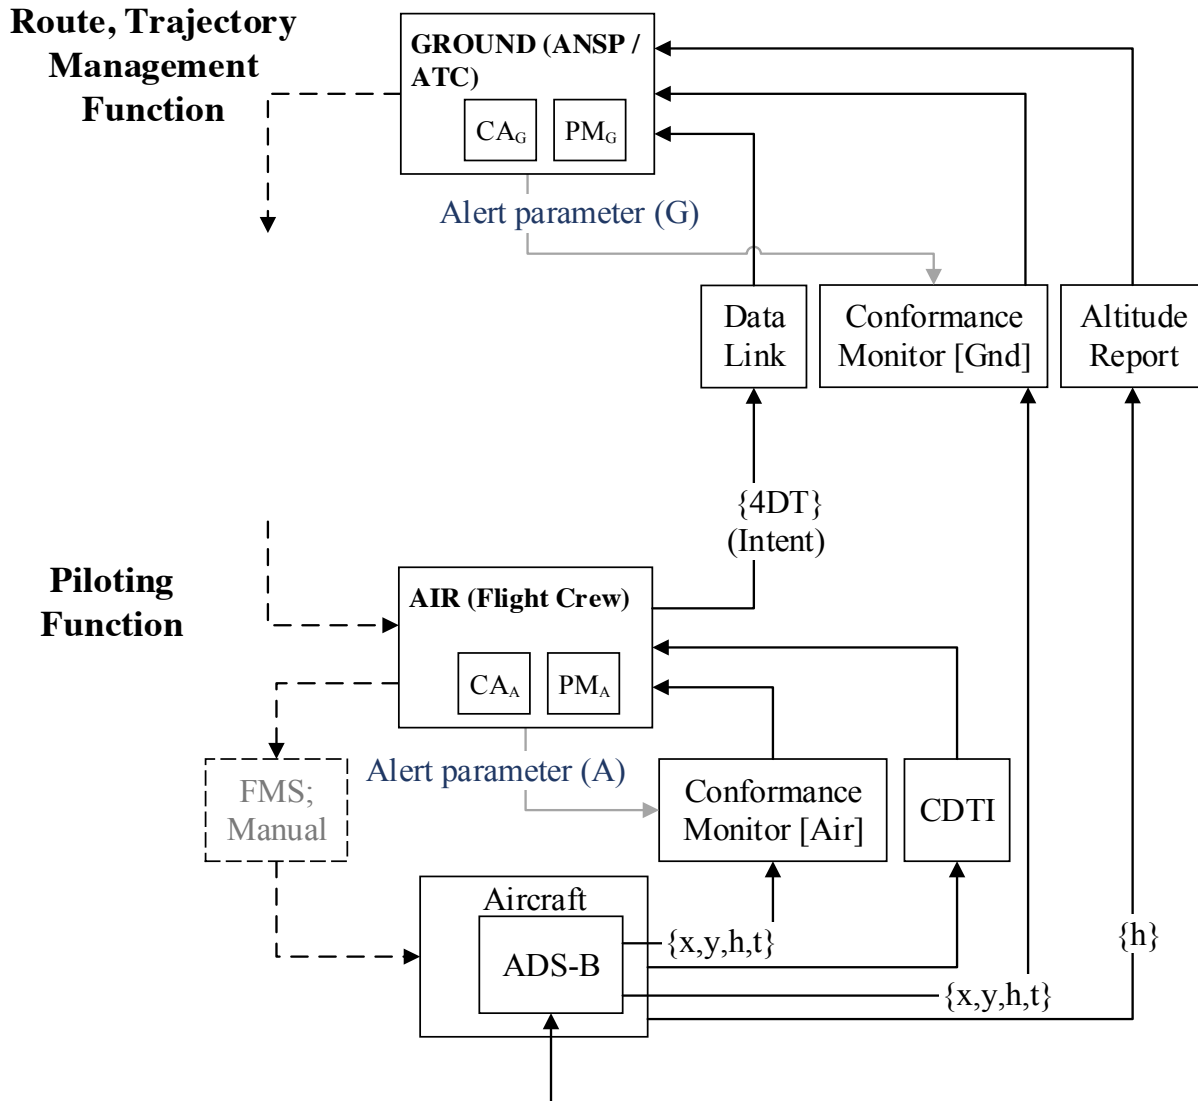
Figure 19. TBO Conformance Monitoring Control Structure

6.3.2   Ground-based Conformance Monitoring Model

Qualitatively, the conformance monitoring consists of a comparison between the measured state of an aircraft in 4 dimensions and its intended state, also in 4 dimensions (three positions and time). If the aircraft deviates from its intended state by more than some tolerance, the conformance monitor will alert the ground controller. The ground-based monitor should have a model of every TBO-enabled aircraft within its jurisdiction.

The TBO ConOps recognizes that navigation performance will never be perfect, so there is a tolerance bound, within which the aircraft can still be considered to "conform". Therefore, in addition to 4D states, the ground-based conformance monitoring model is a function of traffic density, the operation type, and required navigation performance. The tolerance bounds will also be a function of a wake turbulence model, and predicted downstream traffic flows. Finally, the

ground control agent may set an "alert parameter". Table 20 summarizes the content of the model.

Table 20. Conformance Monitoring Model Variables

Conf. Monitoring
- Measured latitude, longitude, altitude for all aircraft
- Intended latitude, longitude, lltitude for all aircraft
- Measured Traffic Density
- Type of Operation (e.g. cruise, descent)
- Navigation Performance (RNP, RTP)
- Wake Turbulence Prediction, Measurement
- Elliptical Conformance Volume
- Downstream, Predicted Traffic Flow
- Alert Parameter

6.3.3 Formal Conformance Monitoring Model

Following is a more rigorous, mathematical formulation that can be used later in the systems engineering process for more detailed analysis, to develop a formal specification, or to generate and integrate into MBSE tools.

The process model for the ground-based control agent (i.e. the ANSP) is simply,

$$\mathbf{PM_G} := \{\mathbf{Conf_i}, \dot{\mathbf{x}}_i\} \tag{43}$$

where $Conf_i$ represents whether each aircraft, i, in the ANSP jurisdiction is conforming with its prescribed trajectory, and $\dot{x}_i$ represents some prediction of the future states of each aircraft. The ground-based control actions are unspecified in the TBO ConOps' section regarding conformance monitoring, but it can still be inferred that the ANSP control algorithm is a function of conformance and that the ANSP will generate an action that enforces conformance. This analysis is for demonstration purposes, and additional analysis of the TBO ConOps reveals other details about ground-based control algorithms, available actions, and models of the airspace. Appendix A contains further analysis regarding ground-based control agents.

Similar to a control agent, conformance monitoring develops a model of the system and contains an algorithm to determine whether certain criteria are met. However, conformance monitoring functions as a sensor and does not have control authority based on what is proposed in the TBO ConOps. The sensor behavior is a mapping from measured variables to controller inputs (see equation. 17). For the ground-based conformance monitor, this mapping is

$$\mathcal{B}_{\mathcal{L}G} := \mathcal{V}_{mG} \rightarrow \mathcal{I}_{G_c}. \tag{44}$$

$\mathcal{I}_{G_c}$, is the signal going to the ground control agent regarding conformance, and the functional behavior of the sensor is

$$\mathcal{V}_{mG} = \mathcal{L}_G \times D_{c,i} \tag{45}$$

$\mathcal{L}_G$ is a model of the airspace state and $D_{c,i}$ is the decision criteria regarding conformance of aircraft i. The "ground", (or ANSP) conformance model is defined as the set of dynamic variables,

$$\mathcal{L}_G := \{z_{int,i}, z_{act,i}, \rho, \tau, P_r, W, E_{cm}, F_D; i \in \mathbb{G}\} \tag{46}$$

Where

$$z_{int,i} := \{G, C, t\}_{int,i}$$
$$z_{act,i} := \{G, C, t\}_{act,i}$$
$$\rho := \textbf{Traffic Density}$$
$$\tau := \textbf{Operation Type}$$
$$P_r := \{\textbf{RNP}, \textbf{RTP}\}$$
$$W := \textbf{Wake turbulence model}$$
$$E_{cm} := \textbf{Elliptical conformance model}$$
$$F_D := \{F, z_{int,i}\}$$

and G is the ground track, C is the climb performance, and t is time of arrival, constituting the aircraft state, z. The int and act subscripts represent intended and actual performance, respectively, for the i^{th} aircraft within ANSP jurisdiction, which is the set $\mathbb{G}$. RNP and RTP represent standard definitions of navigation and time performance, and F_D is a downstream flow model consisting of a general flow model (F) of a particular airspace and a prediction of aircraft arrivals into that space.

Criteria for determining whether an aircraft conforms with its assigned trajectory are,

$$D_{c,i} := \left\{ z_{act,i} \middle| z_{act,i} \notin \bar{\bar{z}}_i \left(z_{int,i}, E_{cm}, a_G \right), \forall i \in \mathbb{G} \right\} \tag{47}$$

where $\bar{\bar{z}}_i$ is an allowed volume for aircraft i, as a function of the intended aircraft state in time, the elliptical conformance model at a given time (E_{cm}), an alert parameter set by the operator (a_G), and $\mathbb{G}$ is the set of aircraft under ANSP jurisdiction. Evaluation of Equation (50) to True indicates that aircraft i does not conform to its assigned trajectory. This is a relatively primitive form of conformance monitoring that simply alerts the user—either a ground controller, flight crew, and/or other software functions that require information about conformance—whether the aircraft is following the assigned trajectory. More advanced forms of conformance monitoring could be developed, but the definition in equation (50) is consistent with the TBO ConOps.

Returning to the ground-based control agent (ANSP) in Figure 19, its control algorithm consists of—at a minimum—use of some decision on conformance alerting. That is,

$$\textbf{CA}_G \supseteq \textbf{Conf}_i \tag{48}$$

where Conf_i is obtained, according to equation 9, via signal processing of the input signal $\mathcal{I}_{G_c}$.

6.3.3.1   Airborne Conformance Monitoring Model

The qualitative description of the airborne model is similar to the ground-based model. The significant differences are that the airborne monitor is responsible only for monitoring its own conformance, as opposed to the entire airspace. Conformance alerting is a function of flight crew input, rather than air traffic controller input. See Appendix A.1, Airborne Conformance Monitoring Model for a similar, formal model of the airborne component.

6.3.3.2   Summary of Conformance Monitoring Models

In addition to the differences between ground and airborne models, conformance monitoring also assumes that other sources of information such as traffic density, operation type, performance requirements, and other factors are derived from the same sources as the ground-based conformance monitor. Further refinement of the process model variables reveals other common sources of information. For example, the TBO ConOps and NextGen in general prescribe ADS-B (Automatic Dependent Surveillance—Broadcast) and GNSS (Global Navigation Satellite System) as the primary sources of aircraft state and intent information for both ground and airborne avionics.

Before performing an analysis of the above model, a few observations can be made about the conformance monitor, its models, and its algorithms. With respect to the conformance model, the TBO ConOps provides a relatively detailed accounting of some of the variables, but little about others. For example, it is well documented that both actual and intended aircraft states will be derived primarily from ADS-B surveillance data. Actual position is measured via satellite surveillance (GNSS), but the intended trajectory is ill-defined in the referenced section of the ConOps[21].

Traffic density is an important factor in determining whether to alert the relevant agents about non-conformance. The conformance monitoring concepts lacks a description of how traffic density is calculated, updated, and which agent(s) has responsibility to do so. It is beyond the scope of this research to determine this parameter, but the point here is that this method of modeling provides a rigorous way to catalog all the components, identify the relationships between components, and compare the relative fidelity of definition of these components and relationships.

The preceding model development made as few assumptions as possible regarding the allocation of control tasks, mechanisms for information exchange, authority, and other factors. The model development method has used only the information documented explicitly in the TBO ConOps section on conformance monitoring, and this information has been parsed in order to refine the functional control hierarchy presented at the beginning of this section (Figure 14).

The following section describes the analytical process, but that analysis requires a more complete model in order to represent a valid demonstration. Refer to Appendix A, which applies this control-theoretic modeling approach to the larger TBO concept. The recursive application of this approach results in the allocation of roles and hierarchy among the various control agents,

---

[21] Other parts of the ConOps describe how the trajectory—and thus the intended aircraft states—will be negotiated and agreed-to, although the actual definition of a 4DT is left relatively vague. Other references define a 4DT with more detail, e.g. [Ballin et al., 2008; Jackson, 2010].

including further definition of the roles of human operators and automation. The modeling effort in Appendix A, also captures additional information exchanges, available control actions, and data sources beyond conformance monitoring.

Upon completion of a hierarchical control model, the approach proceeds with a systems- and control-theoretic analytical process.

6.4   *Analysis of the Model*
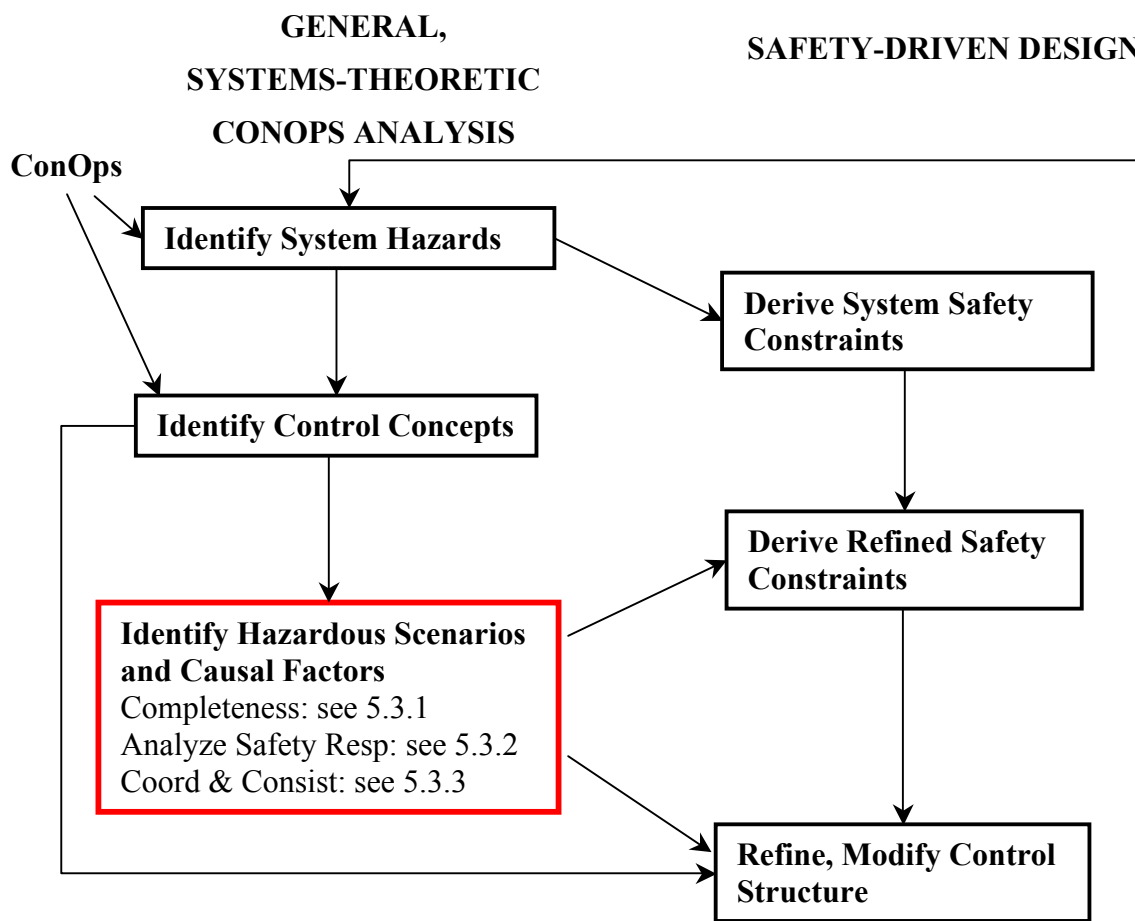
Figure 20. Methodology - Identifying Hazardous Scenarios

This section applies the principles and theory developed in Section 5.3 to the concept of conformance monitoring in TBO. Hazardous scenarios and causal factors are identified according to three general categories—completeness of the control loops, analysis of safety-related responsibilities, and coordination and consistency. Figure 20 depicts this part of the effort, relative to the overall process.

Simply applying these concepts to the model developed above (see Figure 19) would certainly result in a set of potential causal factors. However, the model in Figure 19 is intentionally

incomplete and represents only a subset of elements described by the entire TBO ConOps, and an analysis of such a model could result in the identification of potentially invalid or misleading causal factors. Appendix A documents the more general TBO model, which fills in some of the obvious gaps in the Conformance Monitoring model in the previous section.

The following analysis uses the more general model from Appendix B, TBO Model Development and Analysis (see Figure 40). For the purposes of demonstration, the analysis continues to focus on conformance monitoring but uses the more complete model in order to avoid identifying causal factors that arise simply because of an unfinished modeling effort.

6.4.1 Completeness Criteria for Individual Control Loops

Ignoring the flight operations centers, the model consists of two general control loops. The first model involves ANSP control of the airspace (Figure 21), and the second involves flight deck control of aircraft state (Figure 41). Each loop is analyzed for completeness individually, with the ANSP analysis included below and the flight deck analysis in Appendix A.

6.4.1.1   ANSP (Ground) Control Loop

The model represented in Figure 21 contains more refined control structure information than the relatively abstract, functional model developed in the previous section. Elsewhere in the ConOps, the allocation of control tasks are assigned to various entities—for example trial planning automation, strategic TBO evaluation automation, tactical controller (i.e. human air traffic controller). Humans will be given choices to consider, negotiations to accomplish, agreements to be reached, and thus ATC is the primary control agent, while TBO automation is an input and feedback to the human [e.g. JPDO, 2011, p.41].

Because the ANSP is responsible for all the aircraft in its jurisdiction, individual aircraft are abstracted into a more general "Airspace" component. Upward information arrows in Figure 21 denote this using the i subscript.

89
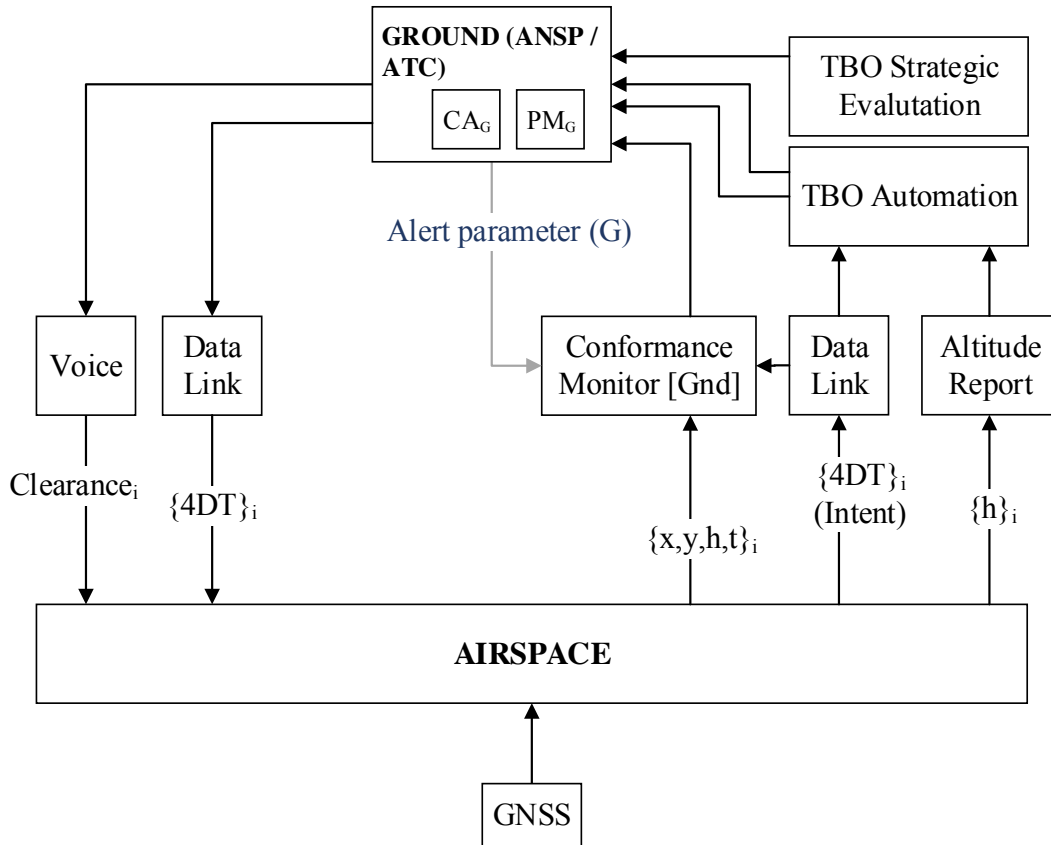
NOTE: i subscripts: X_i represents the control and communication to the i^{th} aircraft

Figure 21. ANSP (Ground) Control Loops

- *Goal Condition*—what are the goal conditions, and how can they violate safety constraints and safety responsibilities?
  The conformance monitoring information in the TBO ConOps alludes to goals of merging, sequencing, and spacing, as well as assuring conformance. Conformance monitoring alone does not describe or specify how these goal conditions are generated or identified; the referenced section does, however, specify the information needed in the algorithm in order to support merging, spacing, and sequencing. See the description of Model Condition and Observability condition below.

  There are discussions and descriptions of merging, sequencing, and spacing elsewhere in the TBO ConOps, but it is left relatively vague in terms of algorithms

or procedures[22]. There is nothing inherently negative about this lack of detail with respect to algorithms and procedures; however, the TBO concept emphasizes the goal of TBO being an improvement over current technologies and procedures for merging, sequencing, and spacing. Thus, the policy for how to achieve this goal must be made explicit. More important for the purposes of this demonstration, however, is the relationship between the goal of merging, sequencing, and spacing and the concept conformance monitoring.

The TBO ConOps implies a policy (and thus a goal) of assuring aircraft conformance in order to meet merging, sequencing, and spacing requirements. Therefore, it can be inferred that the Goal Condition related to conformance monitoring is to ensure conformance. However, the TBO ConOps provides relatively little information about the necessary policies (or goals, or algorithms) for scenarios where aircraft do not conform, to what extent they do not conform, and how non-conformance affects spacing.

In safety-driven design, it is not only important to identify the goal conditions for the various control agents but also to identify how goals can conflict with (or improve) safety. TBO emphasizes conformance and so-called "closed" trajectories. "A closed trajectory is one where the pilot, aircraft automation, the controller, and ground automation all have the same view of what the aircraft is doing, from start of taxi through termination of flight operations" (JPDO, 2011, p.5). This goal is potentially hazardous if adherence to the trajectory will lead to a hazard. Thinking in terms hazards and safety-driven design reveals the underlying assumption that the TBO system[23] will only generate trajectories without conflicts and can identify conflicts when they do occur.

However, an aircraft may have a good reason for not conforming, for example if the aircraft senses loss of separation or that conformance will exceed its capability and ground automation does not have access to this information. The goal of assuring conformance should therefore have the caveat that conformance is only desirable if the 4DT itself does not present a hazard to the aircraft or airspace. This caveat is intuitive and perhaps obvious, but it is either missing from, or obscured in, the TBO ConOps.

- *Action Condition*—how does the controller affect the state of the system? Are the available actions and actuators adequate or appropriate given the process dynamics?

---

[22] Other NextGen improvements that support TBO describe algorithms for merging and spacing, such as interval management and time-based metering.

[23] Based on the TBO ConOps, the responsibility is most likely allocated to automation.

The ANSP controls the airspace using negotiated 4-dimensional trajectories as well as traditional clearances, which are mentioned elsewhere in the TBO ConOps. The "action" is a 4DT delivered via data link or a clearance delivered via voice communication. Are the actuators appropriate for the controlled process?

Another important factor to consider is that 4DT actually implies 6-dimensional action. Aircraft are inertial systems with three dimensions of motion (when airborne), so their state space representations have six degrees of freedom. In traditional air traffic management, the six aircraft states are modified using speed controls, altitude modifications, and heading modification (i.e. vectors). Traditional clearances represent a relatively direct control input to the six aircraft degrees of freedom, while a 4DT represents a mapping to that six-DOF space. That is, 4-dimensional trajectories (may) require an extra step of mathematical or cognitive manipulation in order to obtain speed, heading, or altitude controls.

The Action Condition is thus satisfied with the delivery or acceptance of a 4DT by the ANSP. The Action Condition is satisfied completely with the assumption that the ANSP can issue a 4DT for every aircraft in its jurisdiction (and that the aircraft can actually receive and execute the 4DT). For mixed equipage, the Action Condition is not fully satisfied in the TBO ConOps because not every aircraft in the airspace can be issued a 4DT. Likewise, in delegated airspace, the ANSP does not issue commands to self-separating aircraft under nominal conditions.

In terms of the conformance monitoring, the action condition is satisfied by the generation of a new, closed 4DT based on the current (non-conforming) aircraft state or by a traditional clearance that causes the aircraft to return to its original 4DT parameters.

- *Model Condition*—what states of the process must the controller ascertain? How are those states related or coupled dynamically? How does the process evolve?

  As described in the Action Condition analysis, the most basic states that the ANSP must understand in order to control the airspace are each aircraft's position and velocity. The TBO ConOps section on conformance monitoring specifies several aspects of the conformance model, for example tracked versus intended aircraft state and additional information necessary to make predictions about intended state. What is not clear, in part because the concept is ill-defined with respect to human operators, is how a model—in this case a "mental model"—of a 4DT can be supported by a human agent.

  The ConOps implies a binary output from the conformance monitor, for example a Yes/No conformance alert. This approach only supports a primitive model of conformance. This binary type of model, and the type of sensing that results, does not align with the action and goal conditions described above. That is, the control agent is expected to issue a 4DT or clearance, as a function of some policy or algorithm, but conformance feedback (i.e. Yes/No, see also Observability Condition

92

below) alone does not provide the information necessary to carry out an action that spans multiple dimensions.

Clearly TBO automation will have access—and presumably the ability—to assess current and predicted states, but it is unclear how this approach supports appropriate human behavior. The law of requisite variety [Ashby, 1957] suggests that a binary action condition is only appropriate for processes of two dimensions. Furthermore, a binary model condition can only support an action space of two dimensions.

Based on this analysis, it is unclear whether conformance monitoring satisfies the Model Condition. Because the model should support the execution of a goal and implementation of an action, conformance monitoring should (a) aide in determining whether spacing goals are being met and (b) assist both the TBO automation and the air traffic controller in generating new trajectories or clearances. If the air traffic controller is expected to issue clearances to non-conforming (or even conforming) aircraft, (s)he must be provided information about not only which aircraft do not conform but also the degree to which they do not conform and how it affects separation goals. Perhaps most importantly, and this is a distinguishing feature of human versus computer control agents, an air traffic controller will perform much better if (s)he knows why the aircraft is not conforming [e.g. Rasmussen, 1986; Leveson, 2000b; Dekker, 2013].

- *Observability Condition*—how does the controller ascertain the state of the system? Are the sensors adequate or appropriate given the process dynamics?

  From the perspective of the control agent, conformance monitoring alone does not satisfy the Observability Condition. Given the Goal Condition (Merging, Sequencing, and Spacing) and Action Condition (4D Trajectories, Traditional clearances), more is required than simply an alert about non-conformance. TBO automation will (presumably) have access to the process variables that underlie conformance monitoring, and thus the observability is satisfied by direct access to current and intended aircraft state information.

  Given the fact that ATC will still be able to issue traditional clearances, it may be assumed that the controllers will have access to information similar to that used in current and past operations. It is unclear how existing tools will help air traffic controllers manage 4D trajectories, however. Thus, conformance monitoring does not satisfy the Observability Condition for human operators or even higher level decision support tools or automation that do not have direct access to aircraft state and intent.

  The TBO ConOps emphasizes the importance of achieving the contracted 4D trajectory, which is the essence of conformance. Conformance monitoring represents one important feedback mechanism for determining if an aircraft is maintaining a "set point".

In control theory, a robust control system is one that can reject a wide range of disturbances and maintain a set point. Non-conformance—the inability of an aircraft to achieve or maintain a set point in four dimensions—could be due to any number of disturbances, and the type of disturbance will effect the control policy. For example, a medical emergency to one of the passengers (or pilots) should elicit a different response than an unexpected change in wind direction. Again, it is important to know why the aircraft is not conforming, which could have an impact on both the action and goal condition, as well as the model condition.

While the TBO ConOps does not specify the role of human operators and their interface with TBO automation in great detail, it explicitly states in the description of conformance monitoring and elsewhere that human operators will be part of airspace management in the future. In fact, the TBO ConOps states that human operators will be able to issue traditional clearances. Reasons for providing a vector, for example, are because of non-conformance, conflicting trajectories, or the inability of TBO automation to converge on a viable trajectory. The human operator therefore needs more than just a binary "Yes/No" feedback mechanism to make decisions from a near-infinite set of alternatives. Currently air traffic controllers achieve this with radar, progress strips, and other decision support tools. However, one of the goals for TBO is to increase efficiency, and reducing current margins may render existing tools and mental models obsolete.

If the system will allow human operators to give clearances in order to solve problems with non-conformance, along with other issues, then the system must also support their ability to effectively issue these clearances. Although conformance monitoring is not the only tool that air traffic controllers will use for TBO, it is proposed as a primary source of information. The above analysis of the four control conditions calls into question whether this is appropriate. The analysis suggests that ground-based controllers must be made aware of non-conformance, but also the *degree* to which the aircraft does not conform along with information about *why* it is unable to conform. Thinking in terms of process control and completeness assists in reasoning about human roles, interfaces with automation, and possible control structures as the model in Figure 21 is refined.

This section describes the process and questions used to determine whether and how the model is complete with respect to the four control conditions. More specific scenarios and associated requirements will be presented in the next section. Appendix A provides additional details about causal factors and control structures related to these factors and other aspects of TBO beyond conformance monitoring. The Appendix also includes the analysis of the airborne control loop.

6.4.2 Analyzing Safety-Related Responsibilities

For every hazard there must exist at least one control agent that has responsibility for enforcing the related safety constraint. The control-theoretic corollary is as follows: for every hazardous state variable, there must exist (1) a control policy, (2) an affordance or available set of actions that can regulate the state, and (3) an observer of that state. The above analysis focused on completeness of the control loops, and now the analysis focuses on how the control components enforce (or do not enforce) safety-related constraints.

The TBO Concept of Operations emphasizes the importance of conformance and assigns responsibility to both ground and airborne elements in enforcing this constraint. But how do these responsibilities actually relate to enforcing safety-related constraints, i.e. the prevention or mitigation of hazards? How does conformance monitoring support or detract from these responsibilities and hazards? What are the roles of conformance monitoring in enforcing safety constraints?

The following analysis considers only one of the larger set of system hazards, while the appendix includes additional hazards.

6.4.2.1   Hazard - Aircraft violate minimum separation

Generally, loss of separation occurs whenever the protected airspace of any two aircraft overlap. Traditionally this has been done via a "hockey puck" model, where the protected airspace is represented by a cylinder with 5NM diameter and height of 1000' with the aircraft at its center[24]. Loss of separation occurs when any two of these virtual cylinders intersects. The TBO ConOps proposes other models of protected airspace, which will be discussed shortly.

In safety-driven design, there must be at least one control entity that is responsible for assuring that this loss of separation hazard does not occur (see equation 34). The goal of air traffic controllers, then, in traditional airspace is to generate clearances such that separation minima are always maintained.

One of the objectives prescribed in the TBO ConOps is to ensure that the aircraft conform to their assigned 4-dimensional trajectories. In addition to assuring separation, in TBO the air traffic controllers have the additional goal of assuring conformance. TBO thus satisfies the first general rule for Analyzing Safety-Related Responsibilities (equation 34). That is, safety responsibility is assigned to at least one control agent for the minimum separation hazard.

The next aspect of analyzing the safety responsibilities involves identifying potential conflicts (equation 35). With respect to conformance monitoring and loss of separation, the system must ensure that the respective goals do not cause conflicts. Does the TBO ConOps guarantee that such a condition does not, or cannot exist?

There is a conflict with safety responsibilities if there exists an action that can simultaneously result in the loss of separation hazard and fulfill the conformance condition. Such an action is possible if there are any aircraft (or any other debris or hazardous situation) in the presence of the intended aircraft trajectory or conformance volume.

Any argument against the previous statement must also make the relatively strong assertion that there will never be any conflict along the protected trajectory of an aircraft. An obvious rebuttal to this argument might go as follows. If aircraft α is not conforming and must "work to close" trajectory, it is equally plausible that aircraft β is in the same situation. It also plausible that aircraft β is now on the very trajectory that α is attempting to regain. There are many other scenarios, such as non-TBO aircraft along the trajectory, inclement weather, aircraft

---

[24] This example is typical of en route operations, but separation minima change in terminal operations.

emergencies, and other factors. The goal of assuring conformance must be tempered by the larger goal of assuring separation, as well as preventing other hazards.

Appendix A.4, Analyzing the Safety-Related Responsibilities presents a more formal analysis associated with the preceding discussion, which extends the mathematical foundation developed in Section 5.

Clearly there are many good reasons for conformance, as well as conformance monitoring, and some of these reasons are related to safety. For example, if the ANSP is planning a crossing trajectory behind the trajectory of another aircraft, it must assure that this latter aircraft does not lag behind its assigned trajectory. However, safety-driven development assumes that worst-case scenarios arise, such as attempting to close a 4D trajectory while another aircraft is on that trajectory. The key is identifying the scenarios and then reasoning about (a) how to prevent them and (b) how to mitigate the scenarios if they occur. The following section includes a set of example requirements related to the scenarios described above.

6.4.3 Coordination and Consistency

Conformance monitoring is intended to assure consistency among the various actors in TBO, including the ANSP, flight crews, and operating centers. Despite its intent, it is actually a source of potential inconsistencies and lack of coordination.

Consider again a conformance model, generalized from the development elsewhere. Conformance monitoring is a mapping from surveillance and other data to a binary (or discrete set of) signal(s) that indicates whether an aircraft conforms to the desired trajectory. Conformance alerting is a function of current surveillance data in four dimensions and desired aircraft state in four dimensions, some allowed tolerance volume, and an "Alert Parameter".

This mathematical formulation of conformance monitoring and alerting (equations 47–50) helps identify issues with coordination and consistency, in at least three ways. First, the mapping is a function of an "Alert Parameter". This parameter is available to any agent with a conformance monitor, ground or airborne, and is thus independent and potentially inconsistent. For example, the ground controller sets an alert parameter for his or her own monitor, while the flight crew independently does so for their monitor.

The TBO ConOps does not describe or specify the rationale for including this function, but it may be assumed that it is to counteract alarm overload or over- and under-sensitivity. Furthermore, the TBO ConOps does not specify what the alert parameter entails with respect to human-computer interface design. The ConOps does refer to existing aircraft functions such as altitude alerts for the airborne monitor, but it is unclear how either the ground or airborne agents will "set" these parameters.

In addition to independence and potential lack of coordination due to the alert function, several question arise. These questions also relate to the Model Condition but become increasingly important when considering the coordination of multiple agents.

- What actually constitutes "non-conformance"?
- Does non-conformance mean that an alert is flagged at any instant the aircraft leaves a (continuously updated) elliptical shape?

- Alternatively, does the monitor take an average over a specified time interval and compare that trajectory to the allowable shape?
- How often is the elliptical shape updated? How do the relevant agents receive these updates?
- Who derives the conformance model?
- Is a deviation in one direction given greater importance than a deviation in another direction?

Consider a case where the aircraft is flying the correct 3D shape but does so at the wrong speed. In Figure 22, the aircraft would be flying along the prescribed line but either fore or aft of the elliptical bubble. Alternatively, the aircraft might fly the correct speed but incorrect 3D path and would thus be somewhere perpendicular to the prescribed line. The aircraft's projection onto the line would lie within the ellipsoid. While these factors may be important for the overall efficiency objectives of TBO, they are not important, in and of themselves, for safety.

To reiterate the concepts presented in Section 5, safety is an emergent property and arises from the interaction among components. Even in Figure 22, non-conformance by either or both of the aircraft has different implications for loss of separation, depending on the nature and degree of the non-conformance.
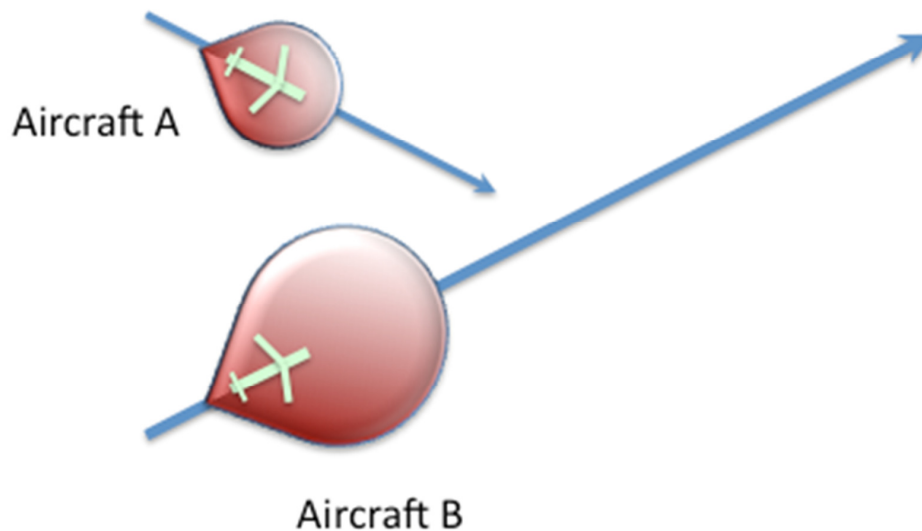
Figure 22. JPDO Proposed Conformance Monitoring Model [JPDO, 2011]

A brief, admittedly oversimplified example is given here for illustrative purposes. Latitude, longitude, and and altitude are measures of length, but their measures are not of equal importance. One pair of aircraft en route are separated in longitude by 0.1NM and vertically by 1NM, while another pair of aircraft that are separated in longitude by 1NM and vertically by 0.1NM. The former pair represents almost zero risk of collision relative to the latter pair.

This example is intended to illustrate the importance of the models that underlie the notion of conformance. The TBO ConOps suggests an elliptical shape of protected airspace, which size, shape and location would be updated at appropriate time intervals to coincide with the intended aircraft operation. Figure 22 illustrates such an elliptical model. What is unclear is which agent

derives and updates the model, or if different agents can do this independently. Is the elliptical model prescribed to each aircraft by the ANSP? If so, how does the ANSP ensure that the flexibility and tolerances comply with the goals and capabilities of the aircraft or its operating centers? Alternatively, do Flight Operating Centers (Airlines) or even the aircraft manufacturers design their own conformance model?

Finally, and perhaps most importantly, the conformance monitor is ultimately used by some control entity to inform its model of the system. As discussed throughout this section, each agent's model is ostensibly used to generate control actions. Two or more control agents, each with different safety-related responsibilities (see Figure 14), have a model of the same process but with different means of updating that process model. Because the various agents have different responsibilities, the different control agents not only form a potentially inconsistent model of the process but also have different means to act on that process.

The next section presents scenarios and requirements associated with consistency and coordination (or lack thereof), along with an option to potentially change the ConOps or pursue a slightly modified control structure.

6.4.4   Summary

The process described above used systems- and control-theoretic techniques to build a model of the TBO Concept of Operations, a crucial aspect of tomorrow's air traffic management system. To identify causal factors in the TBO concept, the model in Figures 4.7 and 4.9 were then queried with respect to completeness of the control components, gaps or conflicts with safety-related responsibility, and coordination and consistency among multiple control agents. These gaps and causal factors then guide the identification of requirements and potential alternative control structures, which is explored in the following section.

As will be shown, this approach shows promise in identifying a different class of problems that existing hazard analysis techniques do not. Section 8 compares these results with those found using existing techniques.

7.   Using STECA in Early Systems Engineering

Hazard analyses or safety assessments should not be used to merely state whether the systems or components are "Safe" or "Unsafe". The results should drive the design of the system, particularly during early concept development when decisions to modify the system or identifying new requirements are most effective and least costly.

The analysis in the previous section is not intended to simply point out potential flaws in the concept. Rather, those results should be used to generate requirements and constraints that eliminate or mitigate against the potential design flaws, while also making undocumented assumptions explicit. The results can also be used to generate and modify the system control structure.

Once the scenarios have been identified, the key to safety-driven design is reasoning about (a) how to prevent the scenarios and (b) how to mitigate the scenarios if they occur.

7.1   *Generating Safety Constraints*

The following requirements are, to the extent possible, solution neutral. For example, requirements related to ANSP control agents do not assume that this function is performed by a particular entity, only that the function must be performed and the information given to the control function (which is also solution neutral in these requirements).

Section 6.3.1 presents the analysis in three groups: Completeness Criteria for Individual Control Loops, Section 6.4.1, Analyzing Safety-Related Responsibilities, Section 6.4.2 and Coordination and Consistency, Section 6.4.3. Requirements are presented in the same order and are based on the relevant sub-section from the previous section, which is depicted in Figure 23.
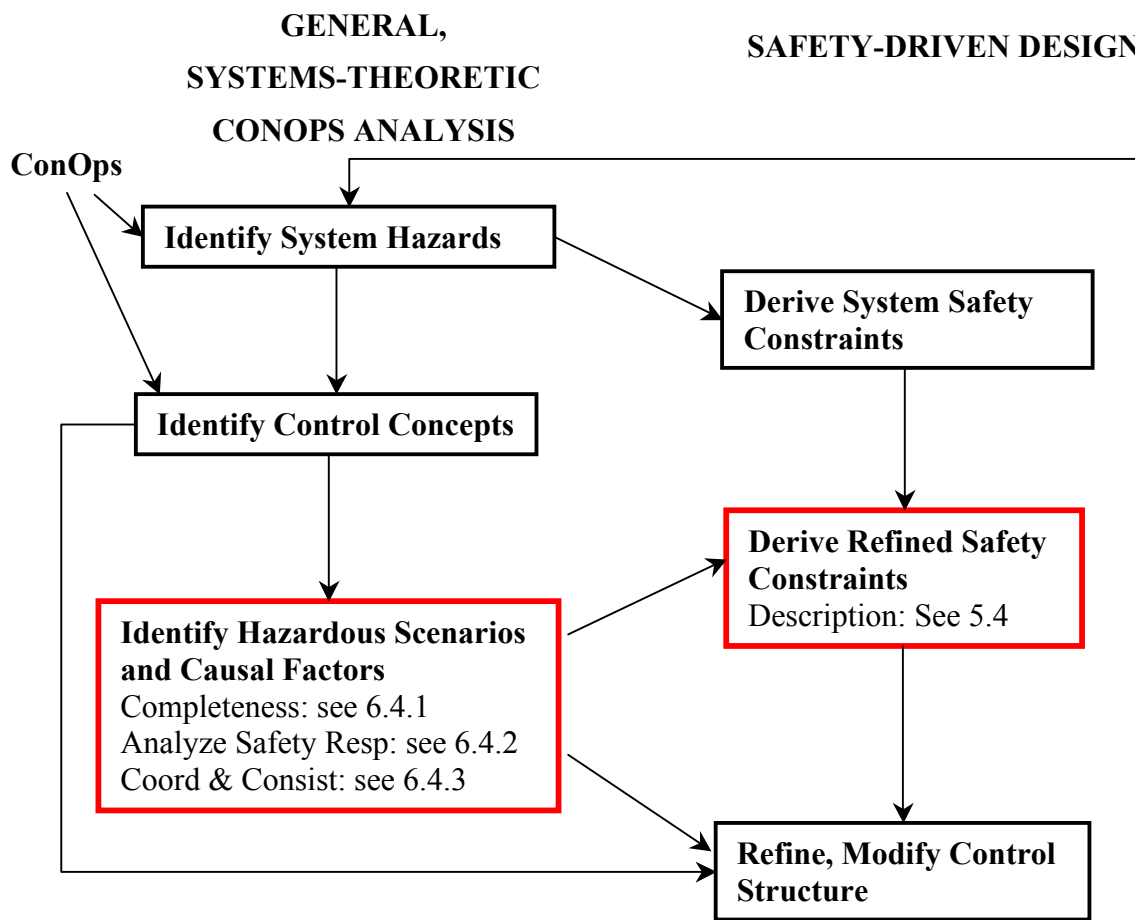
Figure 23. Methodology - Refine Safety Constraints

The example scenarios in Table 21–Table 23 contain pointers to the relevant analysis as well as a reference to the related hazard(s). The scenarios are then followed with example requirements and safety constraints that are based on both the included scenario as well as the analysis presented in the previous section.

Appendix A provides more details about hazardous scenarios and control structures related to these factors and other aspects of TBO beyond conformance monitoring. The Appendix also includes the analysis of the airborne control loop.

7.1.1 Completeness of Control Loops

The previous section describes the process used to determine whether and how the model is complete with respect to the four control conditions. More specific scenarios and associated requirements will now be presented. The example scenario in Table 21 contains a link to the relevant analysis as well as a reference to the related hazard(s). The scenario is then followed with example requirements and safety constraints that are based on both the included scenario as well as the analysis presented above.

Table 21. Requirements Related to "Completeness of Individual Control Loops"

Scenario I.	The control agent, focusing on achieving or maintaining conformance, issues commands that do not achieve the overarching goals of merging, sequencing and spacing. ↑ [H-1] (See Section 6.4.1 - description of Goal, Action, Model, and Observability Condition analysis)
SC-I.1	4D Trajectories must support merging, sequencing, and spacing objectives. ↑ [SC-1] *Rationale*: Self evident based on definition of hazards and top-level safety constraints. *Relevant Causal Factors*: This inadequate control agent may occur if the conformance model is unrelated to the models, procedures, and/or algorithms that must be used to assure separation and achieve merging and sequencing [Goal Condition, Model Condition].
SC-I.1.a.	Trajectories must maintain TBD nmi separation in en route operations
SC-I.1.b.	<other requirements should be levied here for other types of operations>...
SC-I.1.e.	The distance between any two protected conformance zones shall always exceed the required separation minimum. *Rationale*: Aircraft should be allowed to fly at the "edge" of their allowed trajectories, in all directions, and still maintain separation. Sub-requirements here should ensure that the conformance model is updated whenever the separation minima are updated. *Relevant Causal Factors*: The conformance model is not maintained relative to updated separation requirements or across operations (e.g. on descent when separation minima are typically lower) [Model Condition].

SC-I.2.	Current aircraft state, including current 3D position and velocity, must be made available to ANSP control agents. *Rationale*: Current state and immediate intent is a significant (the most significant) information needed to assure separation. *Relevant Causal Factors*: Overemphasis on information about conformance may also override the other types of information needed for merging, sequencing, and spacing (e.g. the location and intent of the aircraft in the airspace) [Model Condition, Observability Condition].
SC-I.2.a.	<Sub-requirements here should ensure that aircraft state information takes precedence over information regarding conformance, particularly when there is a conflict> Requirements related to → Scenario III cover these factors.
SC-I.3	Flight deck must notify ANSP of their intent to deviate from trajectory, including rationale for deviation. This requirement is primarily for tactical maneuvers, which are required within <TBD minutes (or miles) of merging operation or loss of separation threat. *Rationale*: Reasons for deviation help controllers (particularly human operators) in problem-solving. Because the ANSP may be negotiating with several aircraft simultaneous, the rationale for the request may help aide in the decision. *Relevant Causal Factors*: The control agent is unaware of the aircraft's (potentially valid) reasons to avoid the 4DT [Model Condition, Observability Condition].
Scenario II.	Separation is lost within trajectory because aircraft is flying differently than expected in intermediate parts of trajectory. Trajectory definition assures separation at specified 4D waypoints; but waypoints [Goal Condition, Action Condition]
SC-II.1 ...	

7.1.2 Analyzing Safety-Related Responsibilities

Clearly there are many good reasons for conformance, as well as conformance monitoring[25], and some of these reasons are related to safety. For example, if the ANSP is planning a crossing trajectory behind the trajectory of another aircraft, it must assure that this latter aircraft does not lag behind its assigned trajectory. However, safety-driven development assumes that worst-case scenarios arise, such as attempting to close a 4D trajectory while another aircraft is on that

---

[25] Conformance monitoring is described in detail in Chapter 4, as well as in the TBO ConOps [JPDO, 2011].

trajectory. Section 6.3.3 describes some of the rationale and scenarios, while Section 5.3.2 describes the theory underlying the "Analysis of Safety-Related Responsibilities".

Conformance monitoring is only part of a larger set of tools and procedures, and the TBO ConOps recognizes this. Requirements related to the interaction between conformance monitoring and other aspects of TBO, and air traffic management in general, should ensure that separation and basic airmanship take precedence over conformance.

The preceding analysis (Completeness Criteria for Individual Control Loops elsewhere) suggests some additional sources of feedback that may assist in fulfilling these requirements. The requirements and mitigations in Table 22 relate to the role that conformance monitoring has (or does not have) in fulfilling safety responsibilities.

The emphasis in STECA is explicitly considering hazards and safety-related responsibilities when developing and analyzing the role of any component. STECA makes explicit several potentially hazardous assumptions associated with conformance monitoring and its goals. The first hazardous assumption is that it is always desirable for pilots to close their trajectories. A second (perhaps more subtle) issue with the ConOps is that it assumes TBO automation will always generate "safe" trajectories. Further consideration of other hazards such as on-board emergencies, restricted airspace, or any other immediate desire to intentionally neglect following a trajectory reveals the need for other requirements.

Table 22. Requirements Related to "Analyzing Safety-Related Responsibilities"

Scenario III.	ANSP issues command that results in aircraft closing (or maintaining) a 4DT, but that 4DT has a conflict. A conflict in these responsibilities occur when any 4D trajectory has a conflict (conflict could be with another aircraft that is conforming or is non-conforming). See also Scenario I. ↑ [H-1] (See Sec 5.3.2 for a detailed description of this analysis)
SC-III.1	ANSP should not attempt to close the trajectories (i.e. attempt conformance) if a conflict between trajectories exists and updated trajectories cannot be generated within TBD seconds (or TBD NM of separation) ← Scenario I , ↑ [SC-1] *Rationale*: This scenario arises because the ANSP has been assigned the responsibility to assure that aircraft conform with 4D trajectories as well as to assure loss of separation. *See also*: Many of the requirements related to completeness are also relevant here ← Scenario I
SC-III.1.a.	Loss of separation takes precedence over conformance in all TBO procedures, algorithms, and human interfaces *Relevant Causal Factors*: [Goal Condition] For a human operator these requirements could be levied with respect to how the information is displayed; for automation the requirement could be levied in the algorithm in terms of the relative "weight" given to conformance versus generating new clearances.

SC-III.1.a.i.	Loss of separation information must be presented to air traffic controller and/or flight crew *Rationale*: feedback and information should support the goal condition [Observability Condition]
SC-III.1.a.ii.	Loss of separation alert should be displayed more prominently when conformance alert and loss of separation alert occur simultaneously. This requirement could be implemented in the form of aural, visual, or other format(s). *Rationale*: feedback and information should support the goal condition [Observability Condition]
SC-III.1.a.iii.	Flight crew must inform air traffic controller of intent to deviate from 4DT and provide rationale ← [SC-I.3]
SC-III.1.a.iv.	<similar requirements for algorithms; loss of separation should trigger response that takes precedence over conformance>

7.1.3 Coordination and Consistency

Section 6.4.3 describes the systems theory underlying Coordination and Consistency, and Section 6.4.3 then provides a demonstration of how to apply that theory to TBO.

The TBO ConOps prescribes multiple conformance monitors, and these monitors are to be independent. This design decision is consistent with design decisions that often result from traditional hazard analyses - the TBO ConOps implicitly assumes that ensuring independence between component behaviors will lead to a safer design. The next section explores some limitations of traditional hazard analyses, but the point here is that STECA identifies indirect interactions, even when the ConOps assumes independence[26].

Table 23 presents scenarios and associated requirements. The scenarios relate to potential causal factors that actually arise due to these independence assumptions, and the requirements are intended to mitigate against these scenarios by ensuring that the relevant control agents are coordinated. However, modifying the concept might be a better solution, and thus Table 23 concludes with a suggested modification to the structure. The next section develops alternative control structures, using a different example from the TBO ConOps.

Table 23. Requirements Related to "Coordination and Consistency"

---

[26] The existing TBO Safety Assessment further asserts that this independence will serve as a mitigation against certain hazards.

Scenario IV.	ANSP does not issue a DE-conflicting command because it is unaware that the aircraft is not conforming or unaware that the flight crew begins taking action to "close" the trajectory. ↑ [H-1] *Associated Causal Factors:* e.g. ANSP "turns off" or changes Alert Parameter once flight deck has confirmed that it will comply *Alternative Requirements & Control Structures:* Such requirements could be written to eliminate the "Alert Parameter" and require that the black box models of all conformance monitors—every aircraft and on the ground—are identical. Alternative control structures could place the conformance monitoring only in the air (or, perhaps less desirably, only in the flight deck).
SC-IV.1.	Flight deck must notify ANSP of any changes to velocity (change in heading, airspeed, or altitude) *Rationale:* Flight deck typically must request ANSP for a deviation from the filed flight plan, or from the current trajectory. This type of change to the velocity is actually due to the intent of staying on the trajectory, but it changes the aircraft's inertial state (3D velocity); see associated causal factors *Associated Causal Factors:* due to the differences in conformance alerting models, the ANSP may be unaware of the need for change. The ANSP may have a different "Alert Parameter" setting in general, or may have adjusted the setting due to other circumstances (e.g. alarm fatigue, managing other conflicts, etc) Note: This requirement may be levied to either the flight crew or avionics; this design choice would require further analysis of potentially dysfunctional interactions
SC-IV.1.a.	Flight deck must notify ANSP that changes to velocity are due to non-conformance *Rationale:* this is not a "nominal" change, e.g. a change in direction that was part of the flight plan.
SC-IV.1.b.	
Scenario V.	Flight crew does not conform to trajectory, pursuant to the ANSP conformance model. This non-conformance could arise even if the ANSP has instructed the aircraft to do so, and the flight deck has confirmed compliance.

| SC-V.1. | ANSP must issue commands that result in the aircraft closing on the ANSP's own conformance model. That is, the command should directly result in velocity changes that cause the aircraft to enter into desired, protected volume. This clearance is heretofore called a "Close Conformance".

Rationale: ANSP must do more than notify the flight deck that it is not conforming and instruct it to close. This requirement also assumes that the ANSP model of the overall airspace (and thus the conformance model) takes precedence over the flight deck model.

Associated Causal Factors: Flight deck may try to close the trajectory to its own model, or already believe that it is conforming (and thus believe it is complying with the instruction). See also causal factors in ← [SC-IV.1] |
|---|---|
| SC-V.1.a. | ANSP must be able to generate aircraft velocity changes that close the trajectory within TBD minutes (or TBD nmi).

Rationale: TBO ConOps is unclear about how ANSP will help the aircraft work to close trajectory. Refined requirements will deal with providing the ANSP feedback about the extent to which the aircraft does not conform, the direction and time, which can be used to calculate necessary changes. |
| SC-V.1.b. | ANSP-generated clearances used to close trajectories must not exceed aircraft flight envelope ← [SC-V.1.a], ↑ [SC-2] |
| SC-V.1.c. | ANSP must be provided information to monitor the aircraft progress relative to its "Close Conformance" change request

Rationale: See "Associated Causal Factors" listed in ← [SC-V.1] |

7.1.4 Generating the Control Structure

Before proceeding, the definition of a system architecture is necessary. A system architecture is:

> *an abstract description of the entities of a system and the relationships between those entities* [Crawley et al., 2004]*,*

Alternatively, architectures consist of:

> *the structure or structures of the system, which comprise...components, the externally visible properties of those components, and the relationships among them* [Bass et al., 1998].

While there are other definitions of architecture, STECA defines the architecture, or at least part of the architecture, in terms of the hierarchical control structure. The control structure developed in the previous section provides the structure of the system and the relationships between components, like the definition above. These relationships include not only information exchanges, but also hierarchical and authority structures. The control structure can serve as part of a broader definition of the system architecture. Section 5 describes the characteristics of a

control structure, and Figure 24 depicts the inputs and outputs of this aspect of Safety-Driven Design.
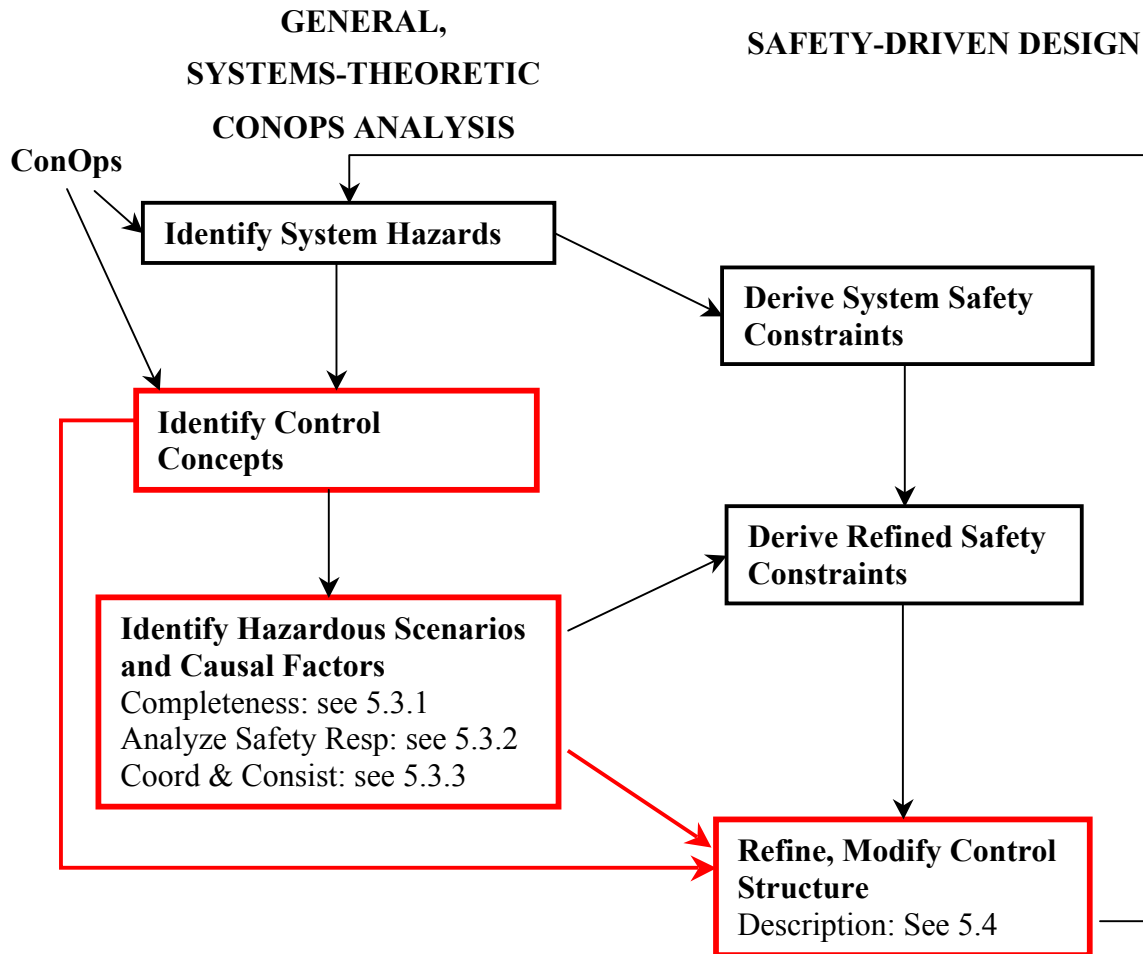
Figure 24. Methodology - Refine Safety Constraints

All of the preceding model development and analysis focuses on conformance monitoring. The requirements (examples shown in Table 21 - Table 23) and control models in Section 6 provide the basis for an control-theoretic description of conformance monitoring. For the purposes of demonstrating how the approach in the previous sections can be used to develop a control structure, consider another example from the TBO ConOps.

In TBO, every approved 4D Trajectory is the result of a negotiation among the ANSP, the Flight Deck, and Flight Operations Centers. Appendix B documents the development of the control model of Trajectory Negotiation, which implements the methods described and demonstrated in previous sections.

Figure 25 depicts the elements of the control model for Trajectory Negotiation, where $\mathcal{K}_y^x \mathcal{K}_y^x$ represents the action from control agent x to controlled process y, and $\mathcal{L}_y^x$ represents the feedback from process y to agent x.

$$A := \text{Air Navigation Service Provider,}$$
$$O := \text{Flight Operations Center,}$$
$$F := \text{FLight Deck, Flight Crew.}$$

The model notionally shows one ANSP, two Flight Operations Centers, and eight aircraft. However, the model (and thus the airspace) may contain as many service providers, operations centers, and aircraft as necessary. Table 24 defines several of the variables in the model, which are omitted from Figure 25 in order to improve visual clarity.

By focusing on coordination and consistency, it can be seen by inspection of Figure 25 that aircraft have the potential of receiving control commands from multiple control agents. These control commands come in the form of approved 4D trajectories. Because the ANSP negotiates simultaneously[27] with flight operation centers and flight decks, even during flight in some scenarios, there is a potential for lack of coordination or consistency. There could be several solutions in the design of control algorithms that could mitigate against these inconsistencies, and requirements could be derived as such.

For example, Figure 26 represents a slightly modified structure that mitigates one of the potential issues with respect to coordination and consistency. A requirement on the control algorithm might read as follows: if the ANSP begins negotiation with Aircraft X, then the ANSP must discontinue (and cannot begin) negotiations with the Flight Operations Center responsible for Aircraft X.

Table 24. TBO Negotiation Structure—Information Exchanges

$\mathcal{K}_O^A$
- Accept 4DT
- Reject 4DT
- Generate and deliver new constraints

$\mathcal{L}_O^A$ —— Request 4DT

$\mathcal{K}_F^A$
- Accept 4DT
- Reject 4DT
- Generate and deliver new constraints

$\mathcal{L}_F^A$ —— Request 4DT

$\mathcal{K}_F^O$ —— Deliver 4DT

$\mathcal{L}_F^O$
- Aircraft Type
- Est. Takeoff Weight
- Range
- Est. & Actual Fuel cost
- Crew availability, allocation
- ...

---
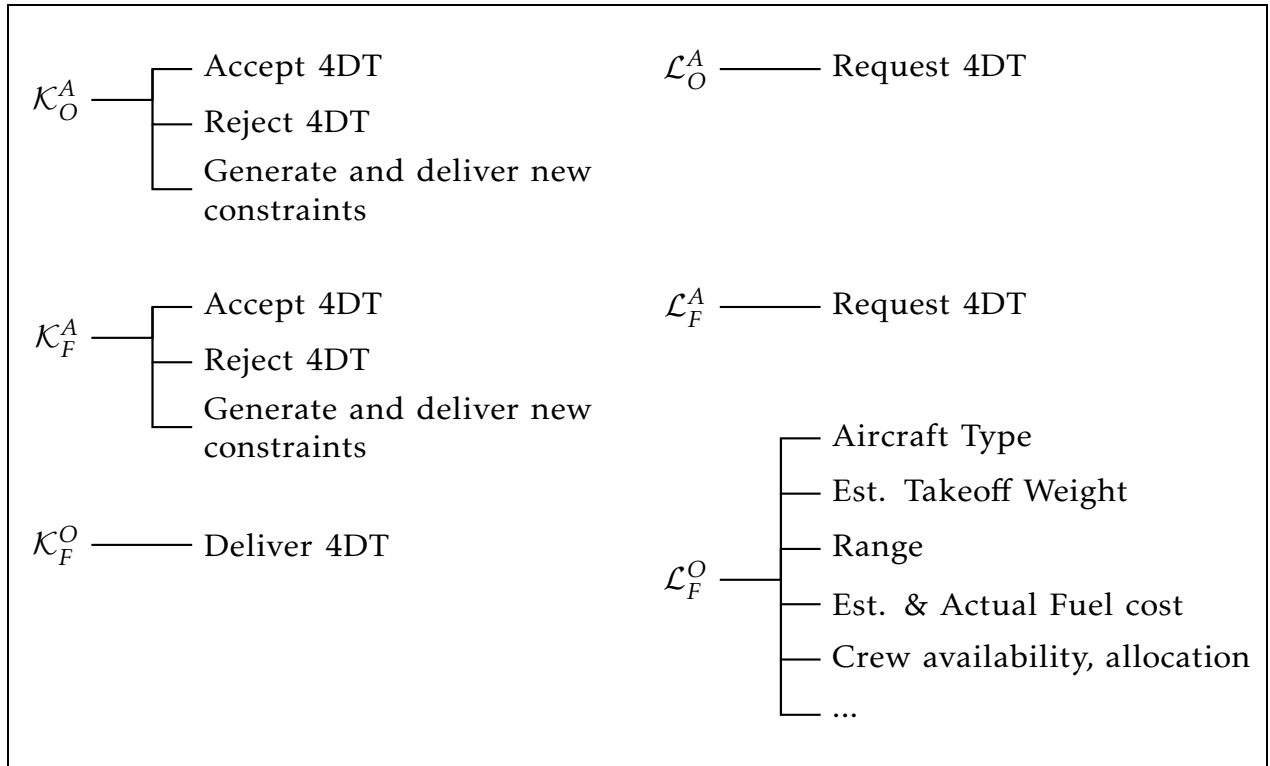
[27] Real-time FOC negotiation with ANSP, i.e. during the aircraft's flight, is in accordance with the TBO ConOps.

However, this problem is perhaps more easily solved with general, control structure modifications. One could implement a high-level requirement that the FOC stops negotiating with the ANSP for all active flights (e.g. within TBD minutes of departure). Such a requirement changes the control structure, from Figure 25 to Figure 27, where the operations center no longer has control authority over active flights and only exchanges relevant aircraft state information. $\mathcal{I}_F^O$ in Figure 27 represents bi-directional information exchanges between operation centers and flight decks, where no control authority exists.
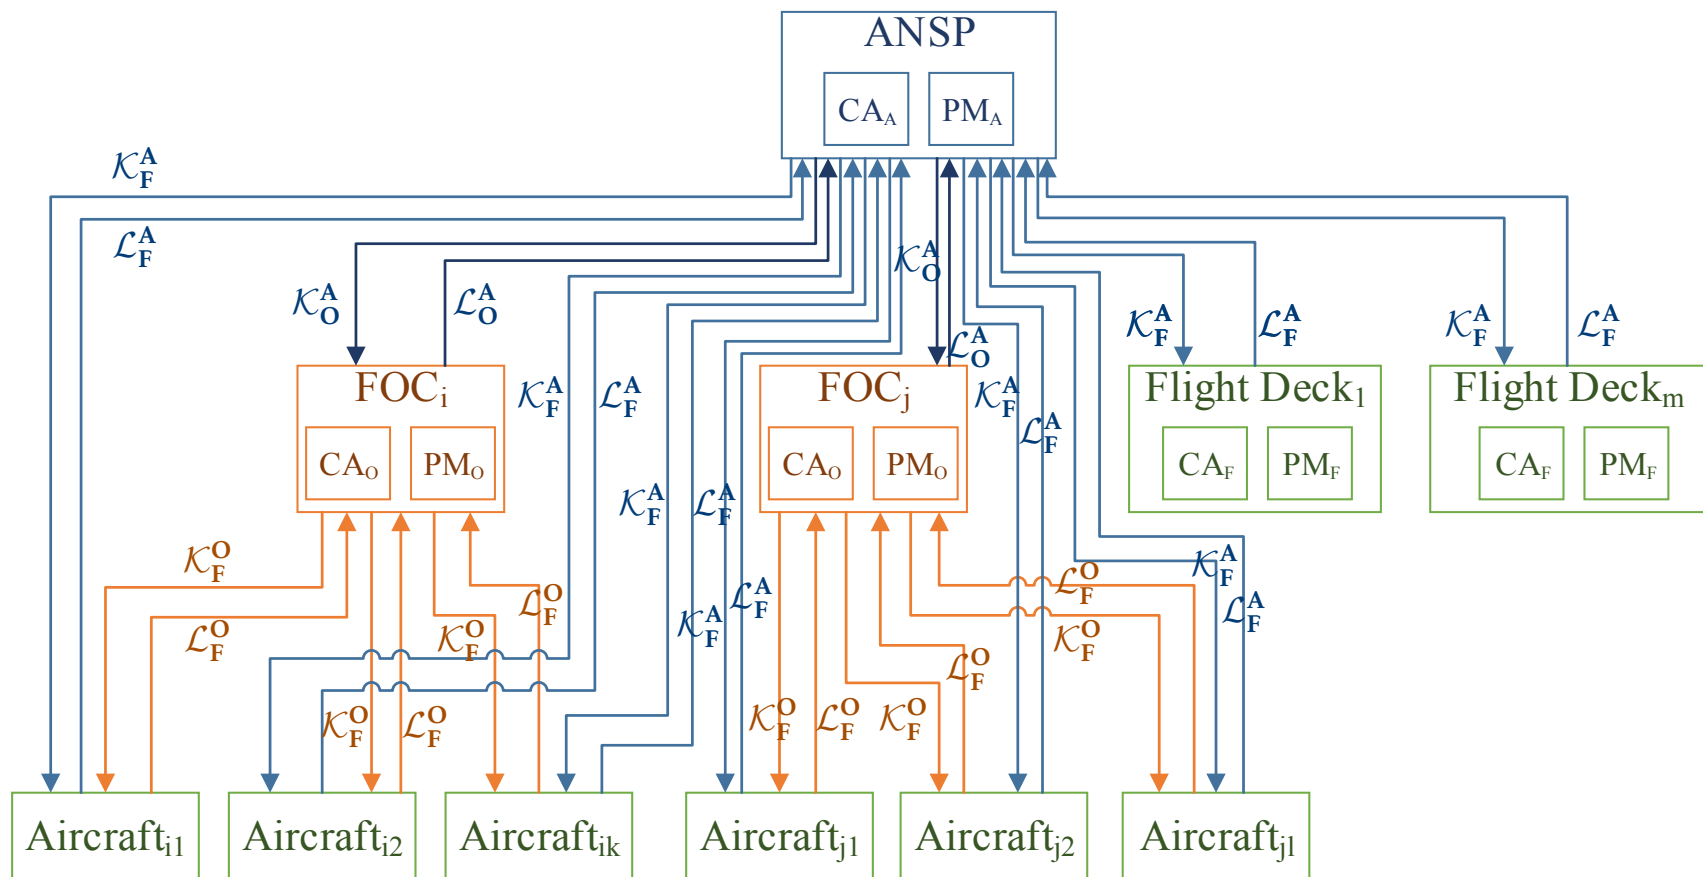
Figure 25. Nominal TBO Control Model - Trajectory Negotiation
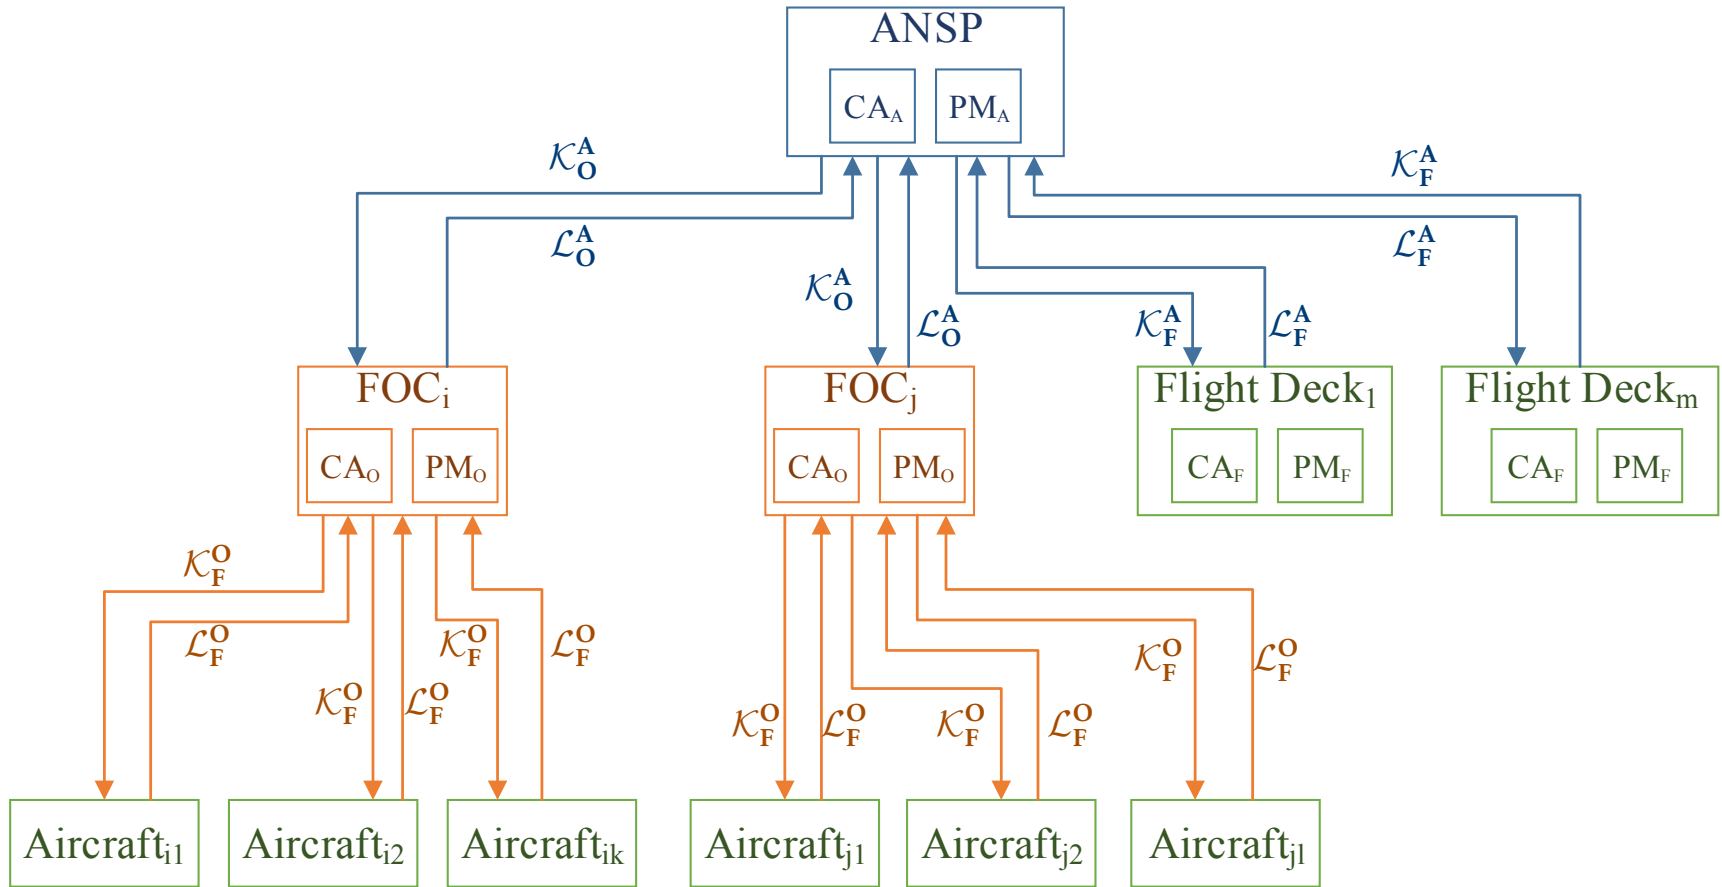
Figure 26. Modified TBO Control Model - Trajectory Negotiation
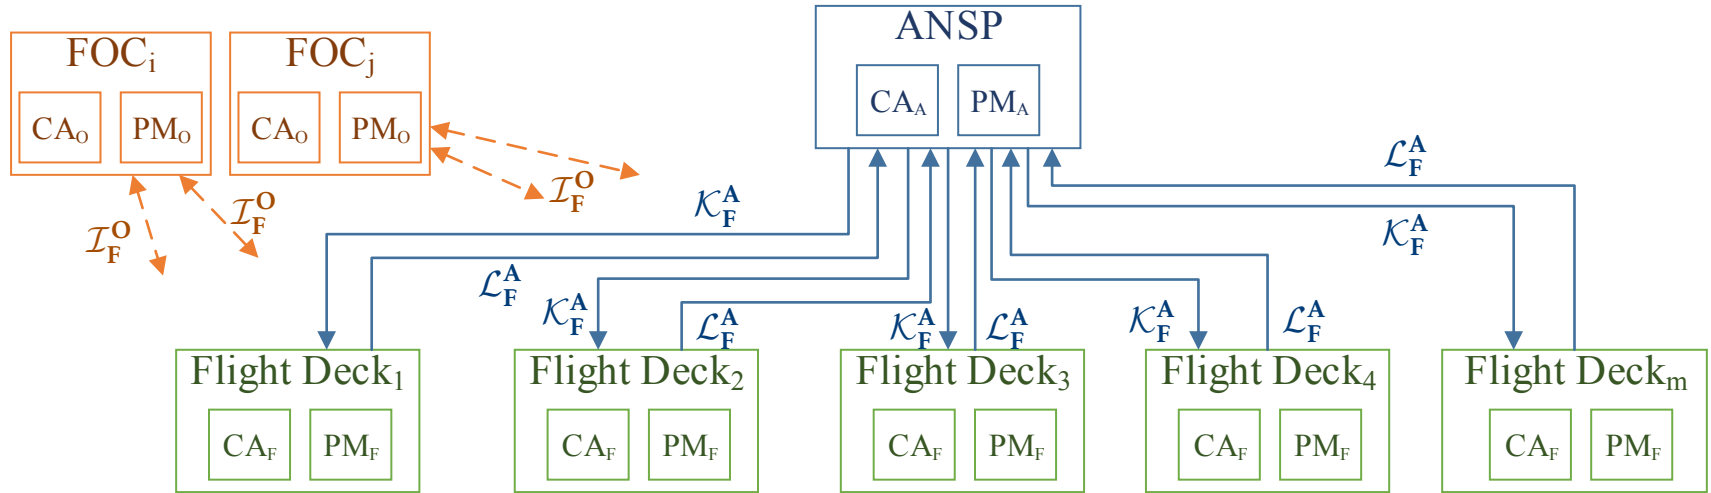
Figure 27. Alternative Control Structure - Trajectory Negotiation

7.1.5 Summary

The process described above uses systems- and control-theoretic techniques to build a model of the TBO Concept of Operations, a crucial aspect of tomorrow's air traffic management system. To identify hazardous scenarios in the TBO concept, the model in Figure 19 and Figure 21 is then queried with respect to completeness of the control components, gaps or conflicts with safety-related responsibility, and coordination and consistency among multiple control agents. These gaps and causal factors then guide in the identification of requirements and potential alternative control structures.

As will be shown, this approach shows promise in identifying a different class of problems that existing hazard analysis techniques do not. The next section compares these results with those found using existing techniques.

8.   Assessment of Results

The assessment of STECA involves comparing traditional hazard analysis techniques with the STECA approach demonstrated in the previous sections. A working group of subject-matter experts applied existing techniques to develop an early safety assessment of TBO, and this assessment provides the basis for comparison.

Though the TBO preliminary hazard analysis provides the basis for assessing this new approach, the following analysis is not intended to be a critique of the TBO PHA working group. Instead, the following analysis asserts that the TBO PHA is representative of general characteristics of a traditional PHA approach and techniques used therein. See Section 3.1 for a description of these general characteristics.

Recall what is necessary for stakeholders to develop a concept. Two significant artifacts of systems engineering, particularly in the early phases, are requirements and the definition of a system architecture. In terms of safety, requirements and architectures should eliminate or mitigate against as complete a set of hazardous scenarios as possible.

Therefore, a successful safety-driven design approach should (1) identify as many valid hazardous scenarios as possible, (2) assist in the identification of requirements and safety-related constraints, and (3) help stakeholders develop a system architecture that eliminates or mitigates hazards. In STECA, the control structure serves as part of the general system architecture definition.

8.1   Existing TBO Analysis - CapSA

The Capability Safety Assessment (CapSA) Team membership included individuals with backgrounds in aviation safety, commercial aviation, air traffic control, government regulations, planning and modeling, and aircraft manufacturing [JPDO, 2012]. Figure 28 shows the approach used by the CapSA team. The following description focuses on the two highlighted boxes in Figure 28 that most closely parallel the approach proposed in this report.

While there are other important steps in the CapSA approach, these steps are only effective if the hazard analysis yields comprehensive results that accurately reflect the types of accident causality found in modern systems. That is, grouping, rating, and sorting hazards and the resulting recommendations (the latter steps in Figure 28) are only as good as the hazardous

scenarios themselves. The following analysis thus focuses on comparing the results of STECA with the hazard identification found in the CapSA report.
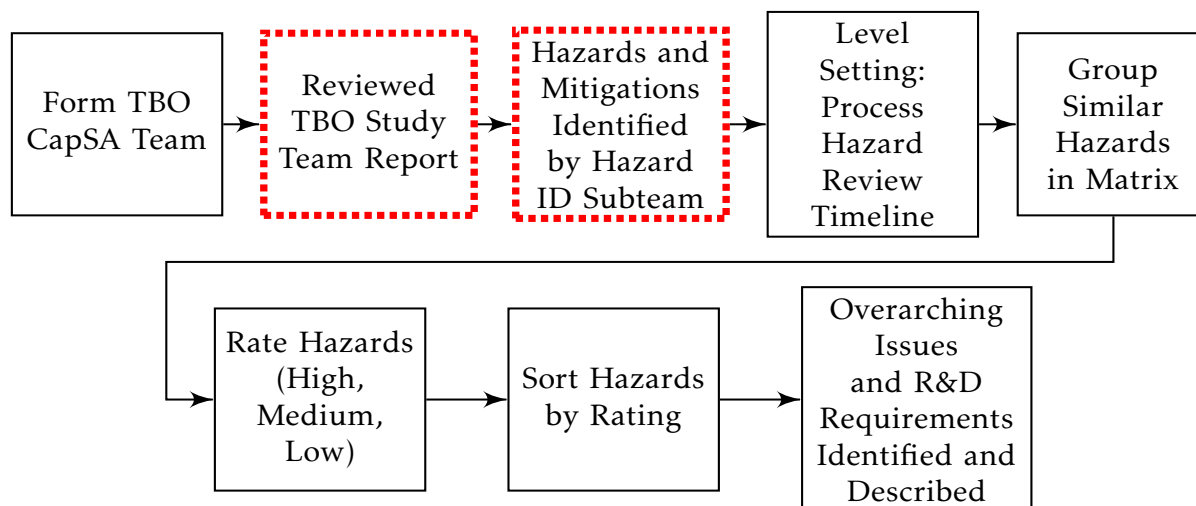
Figure 28. JPDO Safety Assessment Approach [adapted from JPDO, 2012]

The CapSA team and the example analysis presented in this report used the same primary resource document, the TBO Study Team Report[28] [JPDO, 2011]. "The physical, functional, and procedural elements of TBO were extracted and listed from the perspectives of a pilot in an aircraft and an air traffic controller in a ground-based Air Traffic Management (ATM) system" [JPDO, 2012, p.16]. This list formed the basis of the CapSA team's understanding of the TBO ConOps. Where there is insufficient information or detail in the TBO ConOps, the CapSA team achieved consensus on their assumptions and then did a bottom-up failure analysis of each element; each failure is considered a "Base Hazard".

The CapSA team also conducted a top-down analysis, resulting in what is interchangeably called upside down hazard trees or fault trees. The fault trees begin with six "bad outcomes", which are typically called Hazards (and sometimes Undesired or Top Events) in the safety literature [Vesely et al., 1981; Ericson, 2005]. These six events include Loss of Separation, Loss of Control - In Flight, Controlled Flight into Terrain, Runway Collision, Ground Collision, and Wake Turbulence Encounter.

The final step of hazard identification compared the bottom-up failure analysis with the top-down fault tree analysis; duplicates were eliminated and a combined list becomes the final result.

8.2   *Comparison of Existing CapSA to STECA*

---

[28] TheCapSAteamusedVersion1.9.2oftheTBOStudyTeamReportasaprimaryresourcedocument. This thesis used the "Final Report"; there was a 2.0 version—"coordination with members"—before the final release. The documents have minor, cosmetic differences.

Recall from earlier in this section the three criteria for effective safety-driven design[29]: identify a complete set of scenarios and causal factors, derive requirements, and develop a system control structure.

The following analysis is intended to be a one-to-one comparison, to the extent possible. Table 25, Table 26, and Table 27 contain representative artifacts of both the CapSA results and the results from this report.

The first column of each table represents the CapSA results and contains a Hazard Description, Causes (causal factors related to the Hazard Description), and Assumed Mitigations (what the JPDO working group identified to mitigate against the Hazard Description). The second column of each table shows the results using STECA and contain a Scenario (scenarios that lead to a hazard or undesired behavior), Causal Factors (causal factors related to the Scenario), and Requirements (constraints or requirements intended to eliminate or mitigate against the scenario).

8.2.1 Software Behavior

Table 25, column one, presents the risk assessment related to one of the software errors. There are at least two shortcomings in this risk assessment. First, the analysis identifies "design error" as well as "insufficient software testing" as potential causes of a software-related hazard, and the associated fault tree (Figure 29) gives no information about why certain types of design errors might be present. The biggest problem in developing effective software is related to specifying proper requirements. Second, the associated mitigation emphasizes testing. While testing is clearly a worthwhile and necessary activity, it is unclear how testing will reveal a design error—testing is intended to demonstrate that software has implements its specified requirements. Testing also occurs very late in the development cycle when changes can be expensive and ineffective.

---

[29] Guiding the design so that the system is safe and effective is presumably what PHA efforts attempt to achieve.
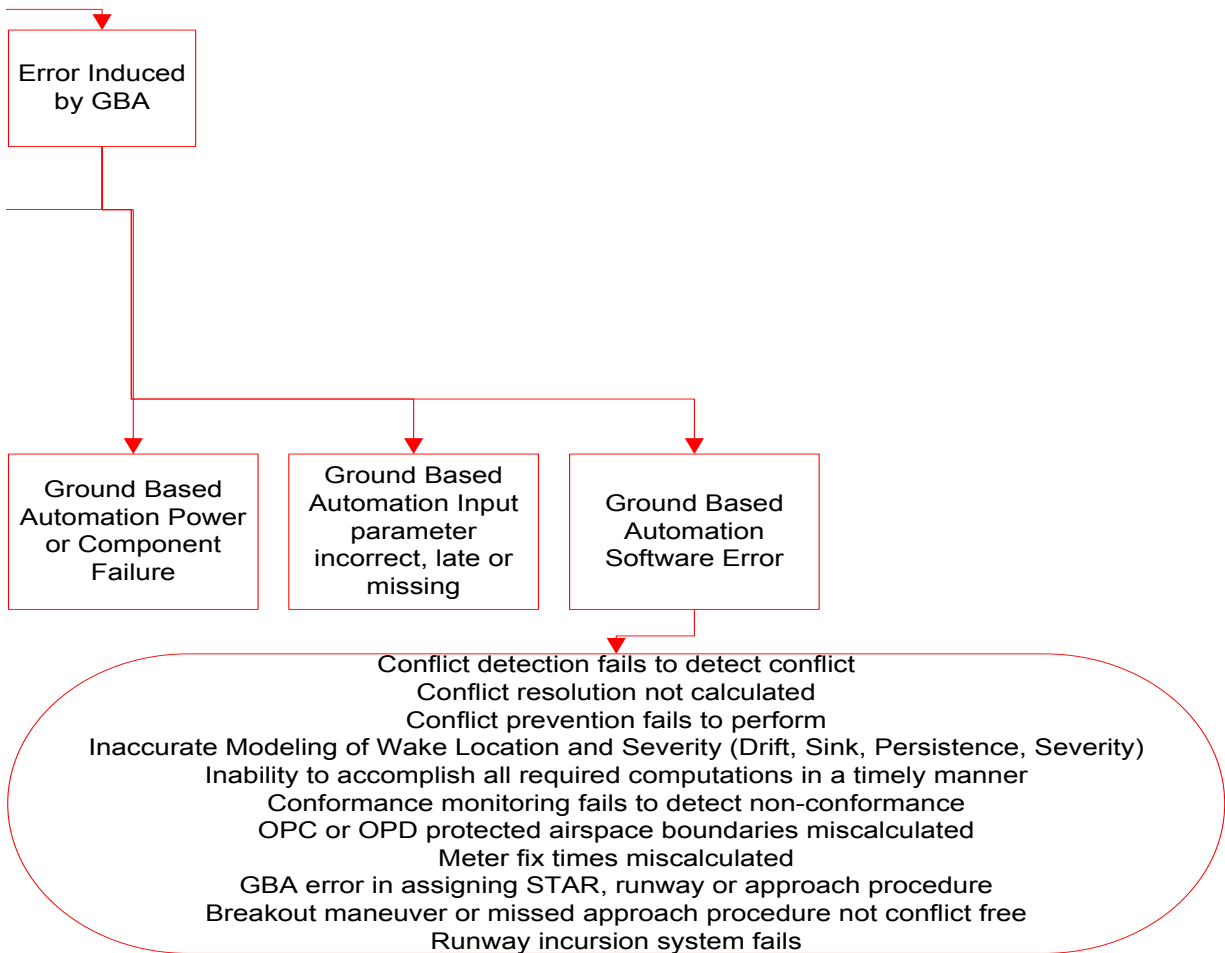
Figure 29. Software in Fault Tree Analysis [JPDO, 2012]

The analysis should help identify potential requirements flaws and guide the design. Because the CapSA analysis focuses on software "failures", it gives no detailed account of flawed requirements, incorrect input parameters, inappropriate behavior, or incorrect algorithm. Software "errors" are only hazardous in the larger context of the system, and the hazard analysis must be able to capture the software's interaction with other components.

The analysis in Section 6 identifies how hazardous scenarios arise with respect to conformance monitoring, and these factors can then be used to derive high level requirements as well as guide the design of the software and the interface with its users. The second column of Table 25 presents one scenario, associated causal factors, and a set of requirements that will drive the software design. Additional requirements are included in the previous section.

As an example, one potential flaw in the design of a conformance monitor is that the conformance model is not updated when the aircraft operation changes or additional aircraft join the airspace. This flaw is an example of an incorrect Model Condition. To mitigate against this flaw, conformance monitoring software must be provided with updated information about separation minima (further requirements would stipulate how and when the separation minima are updated).

Another potential design error identified using the safety-driven approach in this report involves the hazardous interaction among multiple computer systems, which will be explored later in this section (in the sub-section on Component Interaction).

Table 25. Comparison of Software-related Results

Traditional PHA Example [JPDO, 2012]	STECA (this report)
Hazard Description: The software contains an algorithmic or programming error in its implementation. Software that negotiates with ground automation system on 4DT, conformance monitoring, checks for performance capability and navigates to execute the agreed trajectory.	*Scenario*: The conformance monitoring model, i.e. the protected airspace volume, is insufficient or inadequate to maintain spacing
Causes: Design error, coding error, insufficient software testing, software operating system problem; Poor I V&V	*Causal Factors*: This scenario might occur when the 4DT itself has a conflict; The conformance model is not updated to coincide with changing operations (e.g. en route vs. approach); [Model Condition, Observability Condition] The model does not ensure separation because additional traffic has joined the flow and constrained the airspace; [Model Condition, Observability Condition] Different aircraft have different conformance monitors (see "Component Interactions" section below)
Assumed Mitigations: Comprehensive system testing before certification and operational approval. Other ASAS aircraft, if involved, and TCAS. Ground Based Automation would back up in some cases. See and avoid.	*Requirements*: 4D Trajectories must remain conflict-free, to the extent possible Air traffic controllers, flight crews, and/or operations centers must be notified within TBD seconds of an overlap between any two 4D trajectories Conformance volume must be updated within TBD seconds of change in separation minima Conformance monitoring software must be provided with separation minima information ...

8.2.2 Human Operator Behavior

Human error analysis in the CapSA report contains causal factors typically included in fault tree analysis and FMEA. The first column of Table 26 includes one hazard related to the ground control agent (the analysis also includes pilots), and the associated cause is "Human error". There are at least two problems with this cause. While many accidents have been attributed to human error, many behaviors that might be considered an "error" do not result in an accident and can actually be used by the operator to learn and improve his or her behavior [Dekker, 2005].

116

More importantly, like the factors associated with software error, the analysis omits any explanation about why an error occurs and how it might actually lead to a hazard. Because of this lack of definition, the assumed mitigations are equally vague.

The second column of Table 26 identifies hazardous human behavior[30] that may arise due to conflicting goals, missing information, or confusion in the way that information is presented. STECA also leads to specific requirements that can be used to develop the human-computer interface (see the last row of the table). For example, the air traffic controller's responsibility of separating aircraft should take precedence over other goals, which include assuring that aircraft remain on 4D trajectories. One way to enforce this constraint is to ensure that the information presented to controllers enforces their safety-related responsibilities.

While the failure-based approach seeks to overcome human error via automation—Assumed Mitigations in row three, column one of Table 26 assert that conformance monitoring provides a check against the ANSP mistake—STECA takes a different view. For example, STECA seeks to enforce safe human behavior by implementing constraints on the information displayed (ensuring that loss of separation information takes precedence over conformance information) as well as assisting the air traffic controller's understanding of the airspace by requiring additional information that may not be contained in a data-link (flight crews provide intent information when deviating from, or requesting a change in, the 4DT).

Table 26. Comparison of Human Operator-related Results

Traditional PHA Example [JPDO, 2012]	STECA (this report)
Hazard Description: ANSP makes mistake during manual data load into GBA when negotiating a strategic change to the 4DT	*Scenario*: ANSP issues command that results in aircraft closing (or maintaining) a 4DT, but that 4DT has a conflict.
Causes: Human error	*Causal Factors*: This scenario arises because the ANSP has been assigned the responsibility to assure that aircraft conform to 4D trajectories as well as to assure loss of separation. A conflict in these responsibilities occurs when any 4D trajectory has a loss of separation (LOS could be with another aircraft that is conforming or is non-conforming). [Goal Condition] Additional hazards occur when the 4DT encounters inclement weather, exceeds aircraft flight envelope, or aircraft has emergency

---

[30] This sentence intentionally uses the term "hazardous human behavior" instead of "human error"

Traditional PHA Example [JPDO, 2012]	STECA (this report)
Assumed Mitigations: Pilot will have to accept the change; Conformance monitoring; GBA tactical separation; TCAS; Quality of Data check	*Requirements*: Loss of separation takes precedence over conformance in all TBO procedures, algorithms, and human interfaces [Goal Condition] Loss of separation information must be presented to air traffic controller and/or flight crew [Observability Condition] Loss of separation alert should be displayed more prominently when conformance alert and loss of separation alert occur simultaneously. [Observability Condition] This requirement could be implemented in the form of aural, visual, or other format(s). Flight crew must inform air traffic controller of intent to deviate from 4DT and provide rationale [Model Condition]

8.2.3   Component Interactions

Before directly comparing the analytical results, it may be helpful to begin by comparing how interactions are treated in the models[31] themselves. In Figure 30a, interactions are considered explicitly in the control model, and the analysis proceeds by interrogating those interactions in addition to individual component behavior. In Figure 30b, causal scenarios are comprised of "branches" of the fault tree, and for the most part these branches are independent from one another. By using the control model, the analysis in Section 6 showed that different users of conformance monitoring may have conflicting goals with respect to conformance and different models for how the automation behaves. The fault tree analyses in CapSA assume that errors in the air and on the ground occur independently, and the only reason for a breakdown in their interactions is due to communication link failure.

---

[31] Due to space and visualization constraints, Figure 30 includes simplified depictions of both the control structure (a) and fault tree (b). See Figure 22 and Figure 31 for representative depictions of hierarchical control structures and fault trees, respectively.

(a) Control structure   (b) Fault Tree
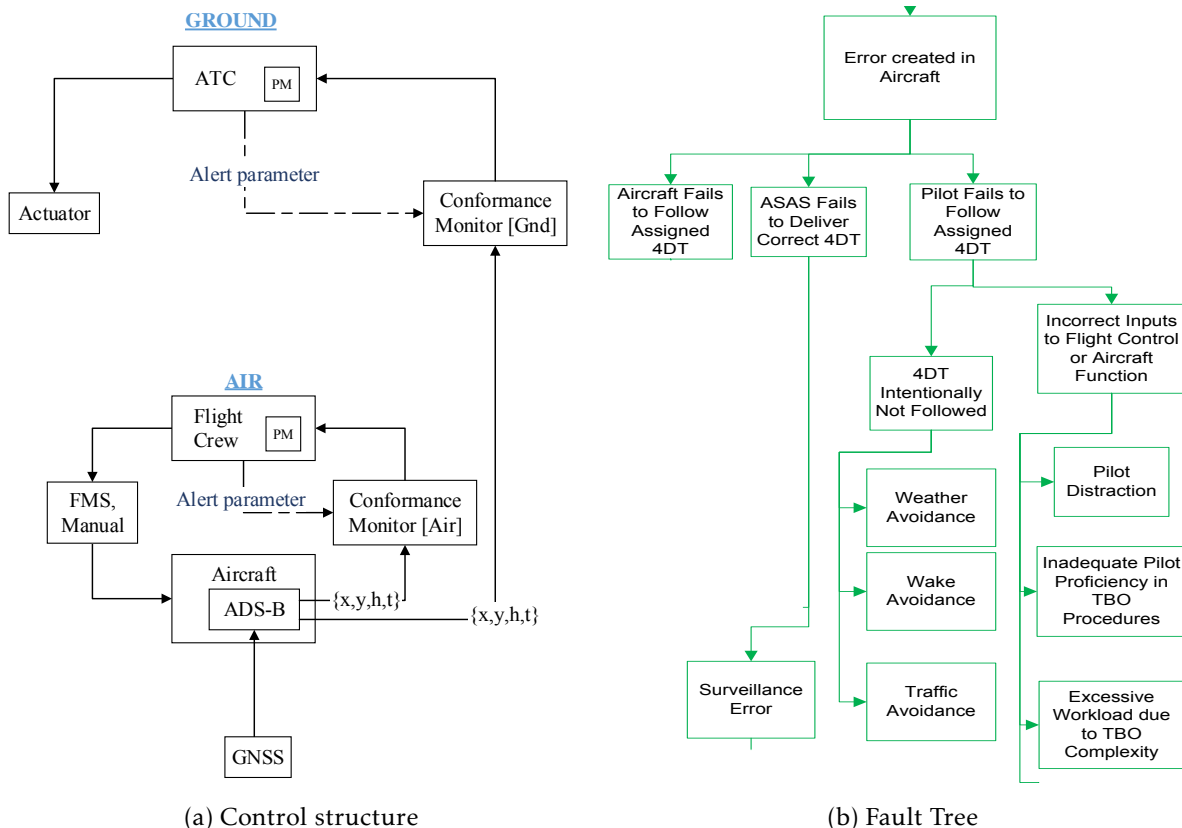
Figure 30. Comparison of Interactions in Analytical Models (see footnote 31)

The Loss of Separation fault tree treats interactions in the following ways. There is a "connection" between automation and human error in which an error in the automation causes human error (e.g. the connection between Input Error and Error Induced by GBA in Figure 31). Though there are other similar connections between events in the CapSA fault trees, most of the causal factors are assumed to be independent.

Additionally, the FTA analysis assumes that the human error will only result from an error in the automation or because of mode confusion. While these are certainly valid concerns, they are not very instructive in terms of the design. At the risk of belaboring the point, there is no information about why the interaction between humans and automation may be confusing, and there are many other human factors issues beyond "mode confusion" [e.g. Rasmussen, 1986; Reason, 2000].

The other CapSA hazards related to component interaction involve communication, but the example in column one of Table 27 again emphasizes component failures (other hazards related to communication include data corruption or interference). The analysis notably omits cases when all components behave nominally but their interactions still result in hazardous behavior.

Table 27 compares specific results side-by-side. STECA finds subtle interactions that go far beyond failures in direct communication. While the independence assumptions in the TBO ConOps are used in the CapSA report as a possible mitigation, STECA shows how this very independence can actually become a source of hazardous behavior. Scenarios in column two of

Table 27 illustrate the potential to lead to dysfunctional or hazardous component interactions, and the associated factors seek to explain how and why such a scenario might arise.

In addition to the requirements generated from this scenario, the previous section described how these results can be used to change the control structure (and thus the architecture and design of the system).

Table 27. Comparison of Component Interaction-related Results

Traditional PHA Example [JPDO, 2012]	STECA (this report)
Hazard Description: The aircraft communications link used to send ADS-B messages fails	*Scenario*: Flight crew does not conform to trajectory, pursuant to the ANSP conformance model. This non-conformance could arise even if the ANSP has instructed the aircraft to do so, and the flight deck has confirmed compliance.
Causes: Equipment failure	*Causal Factors*: Flight deck may try to close the trajectory to its own model, or already believe that it is conforming (and thus believe it is complying with the instruction). [Model Condition] Ground and flight deck have independent Alert Parameters ANSP "turns off" or changes Alert Parameter once flight deck has confirmed that it will comply

Traditional PHA Example [JPDO, 2012]	STECA (this report)
Assumed Mitigations: Redundant equipment; Primary radar; Operational procedures for ADS-B failure	*Requirements*: ANSP must issue commands that result in the aircraft closing on the ANSP's own conformance model. That is, the command should directly result in velocity changes that cause the aircraft to enter into ANSP desired, protected volume. ANSP must be able to generate aircraft velocity changes that close the trajectory within TBD minutes (or TBD nmi). Rationale: TBO ConOps is unclear about how ANSP will help the aircraft work to close trajectory. Refined requirements will deal with providing the ANSP feedback about the extent to which the aircraft does not conform, the direction and time, which can be used to calculate necessary changes. ANSP must be provided information to monitor the aircraft progress relative to its "Close Conformance" change of clearance

8.2.4   Summary

The preceding analysis represents a general comparison of STECA to the approach typically used to perform a preliminary hazard analysis. STECA provides a more complete analysis because it explicitly considers component interactions and uses a more sophisticated model of human and software behavior.

This report has demonstrated how STECA can be used to identify missing and conflicting information in the Concept of Operations and then to use this information to derive requirements and generate a hierarchical control structure. Deriving this engineering information during conceptual design is vital to the successful implementation of tomorrow's complex systems.

LOSS OF SEPARATION (LOS)

Aircraft is following a trajectory that leads to LOS

Error created in Ground Based Automation/ ANSP

Hand-off separation responsibility fails
Transfer of communication fails
Network data communications failures (human – human; automation to human; human to automation; automation to automation)

Error created in Communication between air and ground

Error created in Aircraft

Error Induced by Human Controller

Error Induced by GBA

Spoofed Transmission

Message corrupted in communications equipment

Message not received due to environment

Power or Component Failure

Aircraft Fails to Follow Assigned 4DT

ASAS Fails to Deliver Correct 4DT

Pilot Fails to Follow Assigned 4DT

Inadequate Controller Proficiency in TBO Procedures

Message Corrupted in Ground Equipment

Message Corrupted in Aircraft Equipment

Frequency Congestion

RF Interference

Aircraft Comm Power or Component Failure

Ground Comm Power or Component Failure

Incorrect Inputs to Flight Control or Aircraft Function

Mode Confusion/ Awareness of what GBA is doing in mixed equipage environment

Aircraft avionics Input parameter incorrect or missing

Aircraft TBO Software failure

Performance Limitation

Aircraft Power or Component Failure

4DT Intentionally Not Followed

Controller Distraction

Ground Based Automation Power or Component Failure

Ground Based Automation Input parameter incorrect, late or missing

Ground Based Automation Software Error

Current State of Aircraft Performance Misidentified

Weather Avoidance

Pilot Distraction

Excessive Workload due to TBO Complexity

Wake Avoidance

Inadequate Pilot Proficiency in TBO Procedures

Conflict detection fails to detect conflict
Conflict resolution not calculated
Conflict prevention fails to perform
Inaccurate Modeling of Wake Location and Severity (Drift, Sink, Persistence, Severity)
Inability to accomplish all required computations in a timely manner
Conformance monitoring fails to detect non-conformance
OPC or OPD protected airspace boundaries miscalculated
Meter fix times miscalculated
GBA error in assigning STAR, runway or approach procedure
Breakout maneuver or missed approach procedure not conflict free
Runway incursion system fails

Conflict Detection and Resolution Error (Detect, Decide, Act)

Surveillance Error

Traffic Avoidance

Excessive Workload due to TBO Complexity

Input Error

Aircraft Avionics Input Parameters Incorrect, Late, or Missing

ASAS Software

ASAS Power or Component Failure

Aircraft missing from surveillance
ADS-B contains incorrect position information
Incorrect intent information

Incorrect load of 4DT into Automation

Aircraft missing from surveillance
Failure of hazardous Weather detection for avoidance using 4DT
ADS-B contains incorrect position information
Incorrect intent information
Failure of hazardous Weather detection for avoidance using 4DT

Conflict detection fails to detect conflict
Conflict resolution not calculated
Conflict prevention fails to perform
Conformance monitoring fails to detect nonconformance
Aircraft metering and spacing miscalculated providing incorrect flight guidance
Failure in ability to predict or detect wake hazard
Failure to maintain spatial relationship between VCSPR paired aircraft
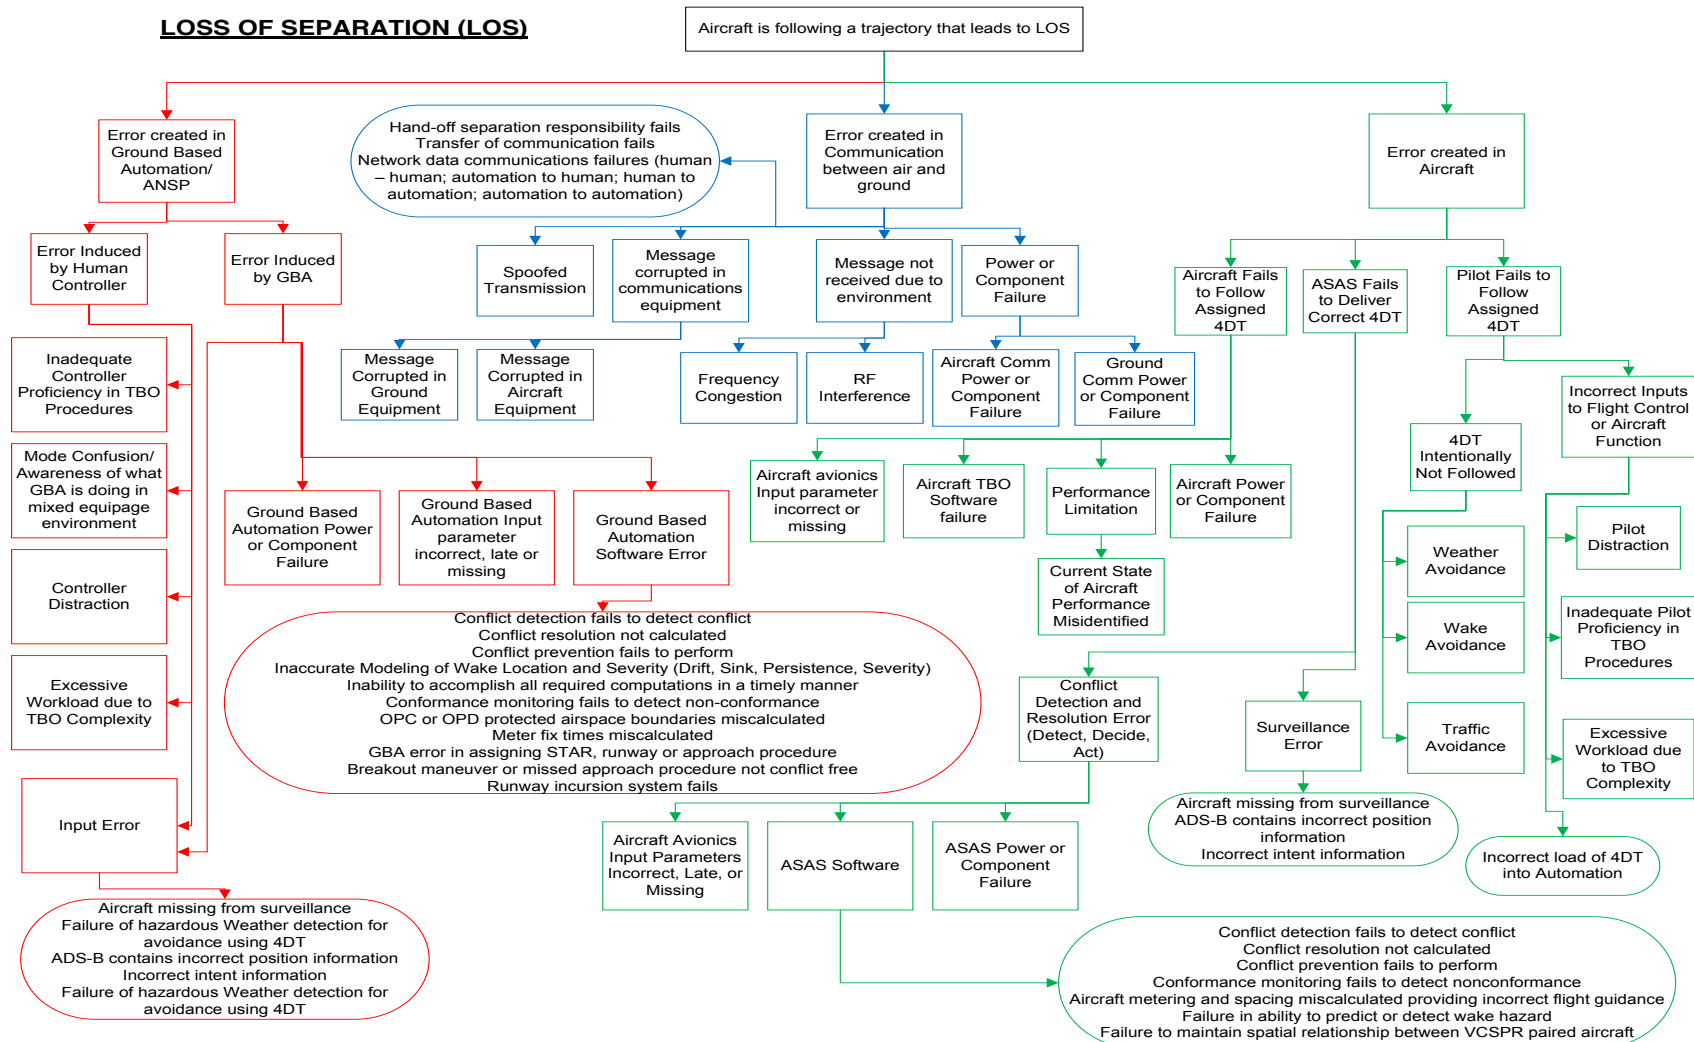
Figure 31. Loss of Separation Fault Tree Analysis [JPDO, 2012]

122

9.   Extending the Human-Controller Methodology

Traditional hazard analysis techniques are grounded in reliability theory and analyze the human controller—if at all—in terms of estimated or calculated probabilities of failure. Characterizing sub-optimal human performance as "human error" offers limited explanation for accidents and is inadequate in improving the safety of human control in complex, automated systems such as today's aerospace systems.

Compared with traditional hazard analysis techniques that stop after assigning human reliability, STPA applies systems thinking to derive both human and automated controller causal factors that relate to flawed feedback, inconsistent process models, and inadequate control algorithms. Using concepts explored in cognitive modeling and ecological situation, this chapter will first offer updates to the human-controller model in STAMP theory in Section 9.1. An update to the human-controller causal-factor analysis will then be pursued in the following section.

9.1   Updating the Human-Controller Model

Figure 38 showcases the general STAMP model for a human controller coordinating with an automated controller to control some physical process. This diagram is a general model—one that can be tailored to analyze any system in STPA—but of particular interest in Figure 32 is the human controller, outlined in the red dashed box.

Starting in the bottom right-hand side of Figure 38, the displays from the automation and the sensors from the controlled process connect to the human controller through any mode of feedback (e.g., visual, auditory, tactile) as designed in the system. The feedback from the displays and sensors then updates the human's mental models through either direct or indirect access as indicated by the solid and dashed lines, respectively. In this diagram, there are two basic mental models the human is aware of: the model of the automation and the model of the controlled physical process itself. These mental models are then used in dynamic (vice static) human control action generation that operates under the influence of written/trained procedures and environmental inputs. Once a control action has been generated, the control action is sent to the automation or the controlled physical process, through controls or actuators, as either a primary or backup control action, as indicated by the solid and dashed lines, respectively. When zooming out in Figure 38 and contrasting the red box of the human controller with the automated controller directly beneath it, note that there is little distinction between the two. The only change to the human controller is an unseen dynamic control algorithm (vice a static, automated algorithm) and an additional process model of the automation itself. Otherwise, the two controllers are identical.
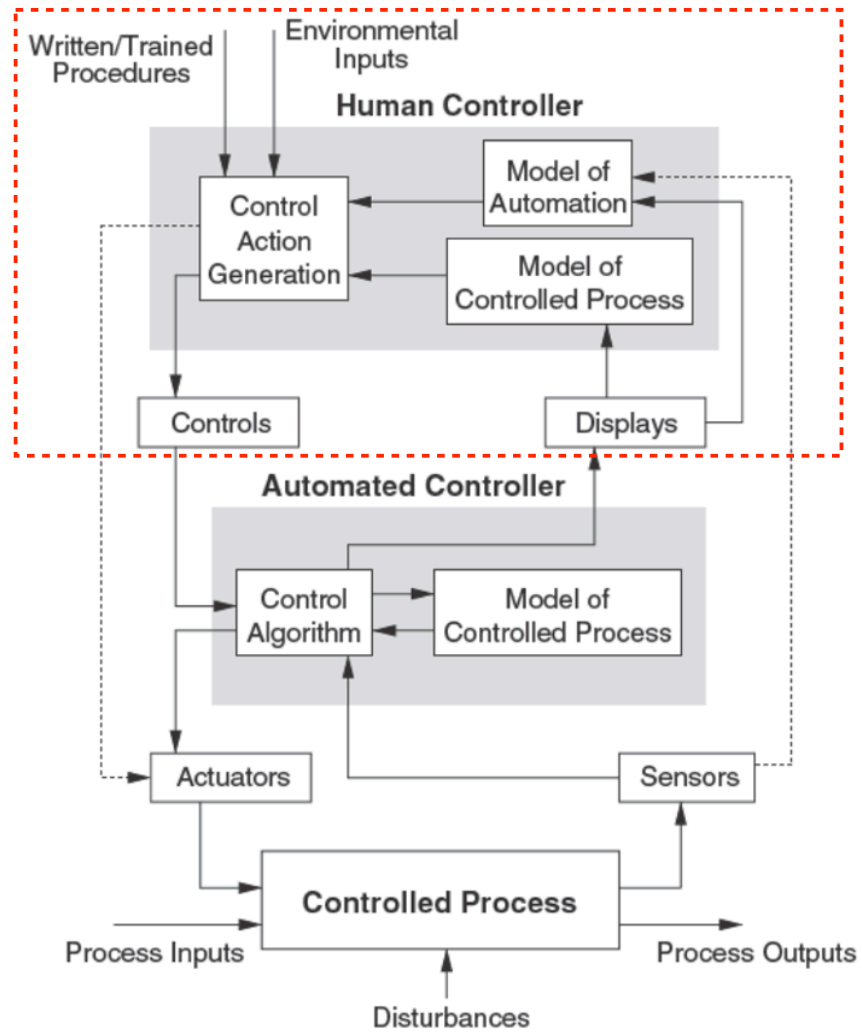
Figure 32. Original Human-Controller Model in STAMP [Leveson 2012[

124

35

In response to this, the model of the human controller can be updated by using the Human Factors concepts explored from in the literature to better reflect traits unique to the human operator [eg. Flach and Voorhorst 2012, Boyd 1995, Rasmussen 1983]. A more comprehensive review of the domain of human factors as it relates to process control is provided elsewhere [Thornberry 2014].

One of the most acclaimed models in basic human cognition was proposed by Air Force Colonel John Boyd and classically known as the OODA (Observe-Orient-Decide-Act) Loop, represented in Figure 33.
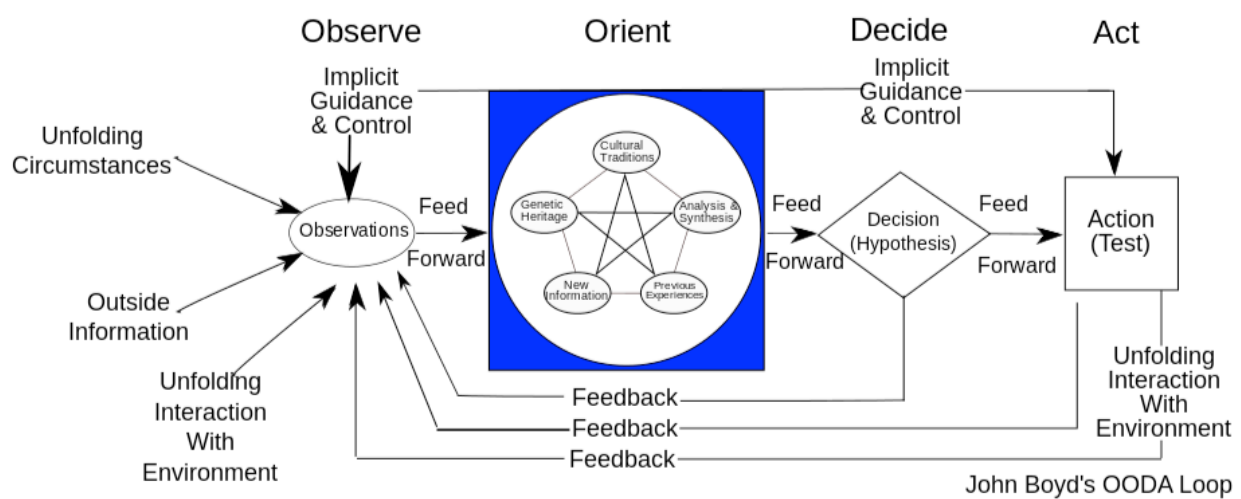
Figure 33. OODA Loop [Boyd 1995]

This model was initially developed to frame and explain both military tactical-level and strategic-level operations, but the ready applicability to other industries and the Human Factors domain quickly made this model popular [17]. As the name suggests, this model has four main components to it: observation, orientation, decision, and action. When performing a task, the conditions and states necessary for the successful execution of that task are first observed. Following this observation, the controller then orients themself to the problem space (work ecology) before deciding upon a course of action and eventually realizing that action. This process is not a set of psychological sequences, but rather a continuous process much like a control loop. One of the greatest strengths to Boyd's model, aside from its simplicity, is that it can be applied across all domains, time scales, as well as to different controllers, whether human, automated, or organizational. The flexibility in Boyd's approach will see basic application to hazard analysis in the following sub-section.

A second model of human cognition that has been used successfully in complex domains is Rasmussen's Skills, Rules, and Knowledge (SRK) Framework as depicted by his model in Figure 34.
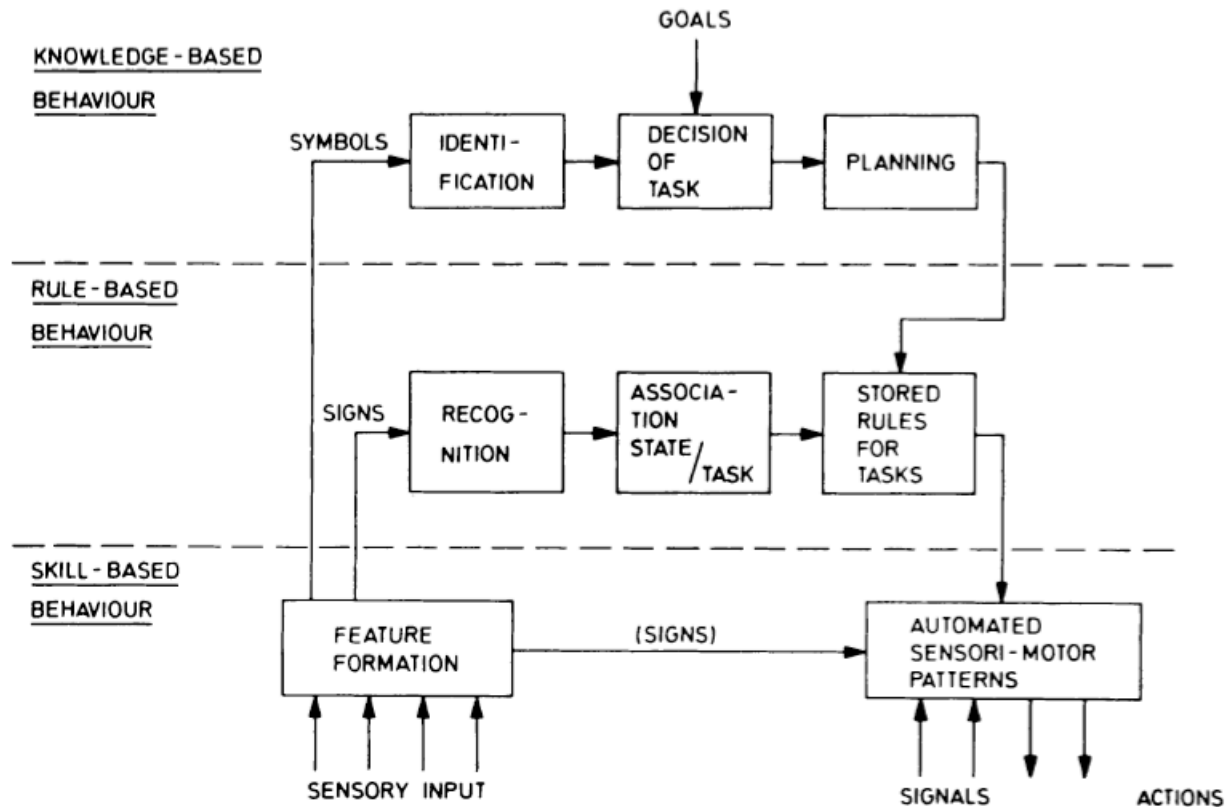
Figure 34. The Skills, Rules, and Knowledge Model [Rasmussen 1983]

Rasmussen's SRK Model is broken down into three separate levels of performance: skill-, rule-, and knowledge-based behavior. Skill-based behavior represents those lower-level sensory-motor skills that take place without conscious thought, like riding a unicycle or dribbling a basketball, for example. Rule-based behavior considers performance that is conducted through the use of previously stored rules within the mind, such as cooking a family recipe, for example, and usually occurs under an implicitly formulated goal. The last category of knowledge-based behavior considers human problem solving with unfamiliar tasks and limited or no rules for control—but with an explicit goal or objective in mind. Learning to play chess for the first time is an ideal example of this.

Of particular interest in Rasmussen's SRK Framework is the way in which he categories feedback for these three behaviors. Skill-based behaviors are fed by a source of feedback known as **signals**—those "continuous quantitative indicators of the space-time behavior of the environment [Rasmussen 1983]." The next level of feedback used to make rule-based decisions is known as a **sign**—physical states or situations in the environment that inform rule-based behavior. Lastly, the categorization of feedback that supports knowledge-based behavior is known as a **symbol**—a functional property in the environment that is generated through human meaning. This idea of meaning, in fact, directly relates back to the previous discussion on the role of feedback in ecological interface design. Given the overall goal of extending the human-controller analysis within the STAMP framework, concepts from both ecological psychology and basic cognitive modeling, to include the SRK Framework, will be applied to the human controller later in this section.

126

Based on these fundamental concepts in human factors, the updates to the STAMP human-controller model is shown in Figure 35.
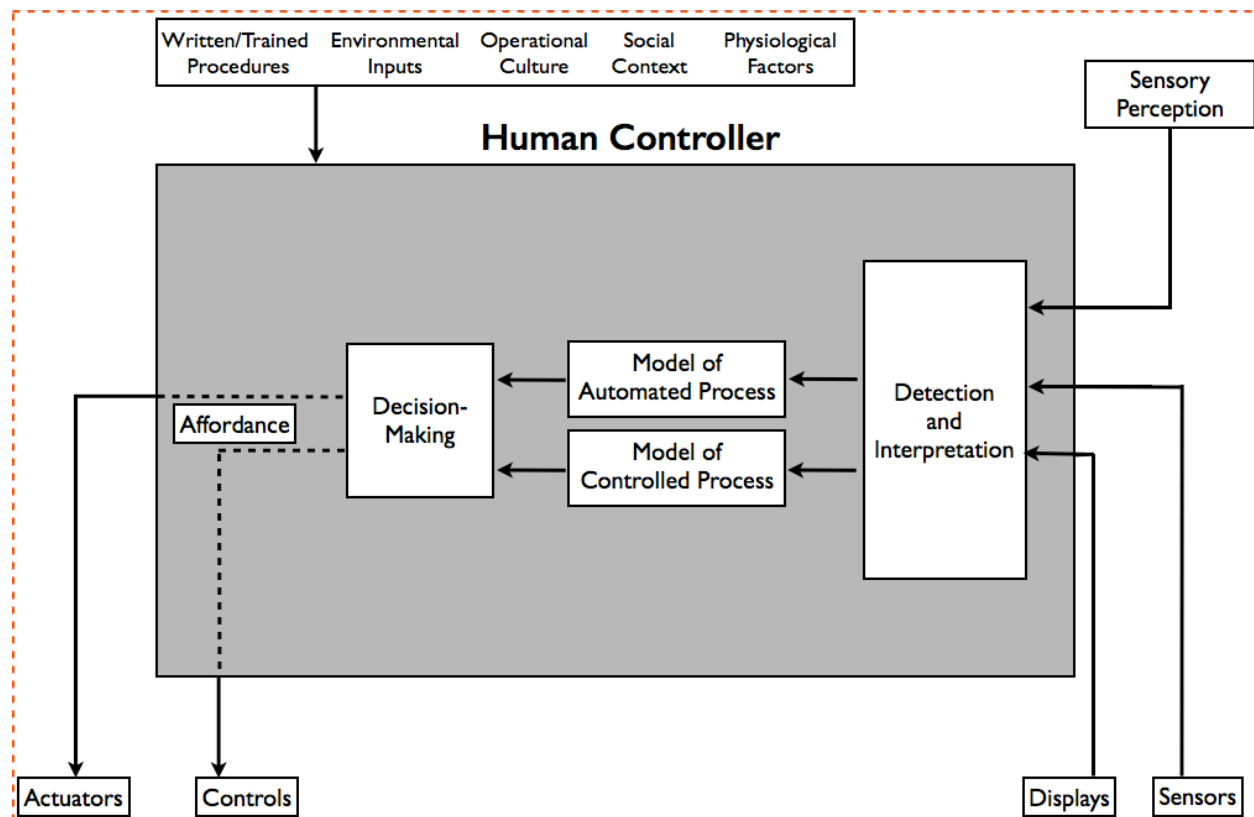
Figure 35. Updated Human-Controller Model

Working from the bottom right-hand corner of Figure 35, the displays from the automation and the sensors from the controlled process connect to the human controller just as they did in **Figure 10**. In addition to these two sources is a clarified tertiary source of the human operator's own sensory perception that was refined from the original "environmental inputs" in Figure 32. This sensory feedback includes that raw visual, proprioceptive, and vestibular feedback the human receives *that has not been designed into the system*, similar to Rasmussen's "signals." The inertial accelerations felt by a human while driving a car or flying an airplane would be an example of this type of feedback.

All three of these feedback categories then feed into a "detection and interpretation" stage through which the mental models are updated. Most generally agreed upon models of human cognition, like the OODA loop and SRK Framework, include some stage of human information recognition or observation, and thus this new category was added to Figure 35. The human operator must accurately process and understand the feedback from the displays, sensors, and

their own sensory perception in order to form the correct mental models. It is important to note that *accurate* human interpretation of the feedback in Figure 35 is defined to be *as the process states exist in reality*. This relates to ecological interface design where the intention is to uncover the situated "state" variables in the ecology and relate them to the human controller through feedback. Therefore, if the process states in STAMP (or STPA, or STECA) that are channeled through feedback are interpreted in any way other than as reflected in reality, the mental models of the human may not align and cause an accident—an allusion to a causal factor that will be explored in the next section.

Following the detection and interpretation of these feedback channels, the mental models of the human operator are then generated or refreshed, just as they were in the original model in Figure 38. From these mental models comes the renamed "decision-making" process for control actions. During operation, the human operator must consider and decide on higher-level goals within the overall control-action decision process. Not only do these higher-level decisions affect lower-level control actions, but these decisions can also transpire across different time scales as well, from seconds to hours—a matter of complexity that for now will be summarized as "decision-making."

Once a control action has been decided upon, the human operator must then afford the action before it can be realized. Although strictly thought of as a property of the environment, affordance (once again) is defined as the coupling between human motor effectivities and the opportunities in the ecology of the system [14]. If the human operator cannot afford an action in time because of a poorly designed interface, for example, or if the human affords a control action unknowingly (e.g., automation surprise), an accident *will* occur when paired with a worst-case set of environmental conditions—more on this idea in the next section.

Once the control action is realized it is then sent to the actuators and/or controls just as in the current human-controller model of Figure 38. The entirety of this process, from detection to control action, takes place under the influence of a variety of external factors as seen in the top box in Figure 35. These include written/trained procedures, other external environmental inputs not related to lower-level sensory perception, an operational culture, a social context, and physiological factors (e.g., stress, fatigue, workload). All of these factors affect how the human controller detects and interprets feedback (e.g. through trained scan patterns) [19], how they assimilate feedback into accurate mental model representations, and how they make decisions accordingly. This is to say that every stage within the entire human operator "box" is affected by these factors and that they must be considered when analyzing the human controller.

Altogether, this updated model discretizes the abstract cognitive mechanisms and affordance component of the human operator when controlling a general system. As seen in Figure 35, these cognitive processes have been refined into the categories of "detection and interpretation," "mental models," and "decision-making" with the addition of "affordance"— the precursor to any control action. These categories are not comprehensive, however, as there are near limitless ways to model human cognition and information processing in the domain of Human Factors. Any one of the cognition components in Figure 35 could also be elaborated further by culling more detail from cognition models and theories. As Dekker warns, however, these detailed models are impossible to prove or disprove [1], and insofar as improving human controller-related hazard analyses, a simple yet precise picture of human cognition is all that is needed. In

pursuit of extending the human-controller methodology within the STAMP framework, this updated model was applied to the causal-factor analysis itself in Section 9.3.

9.2   Updating the Human-Controller Analysis

The original controller causal factors analysis outlined in the red dashed box in Figure 36 focuses on four distinct categories of causal factors.
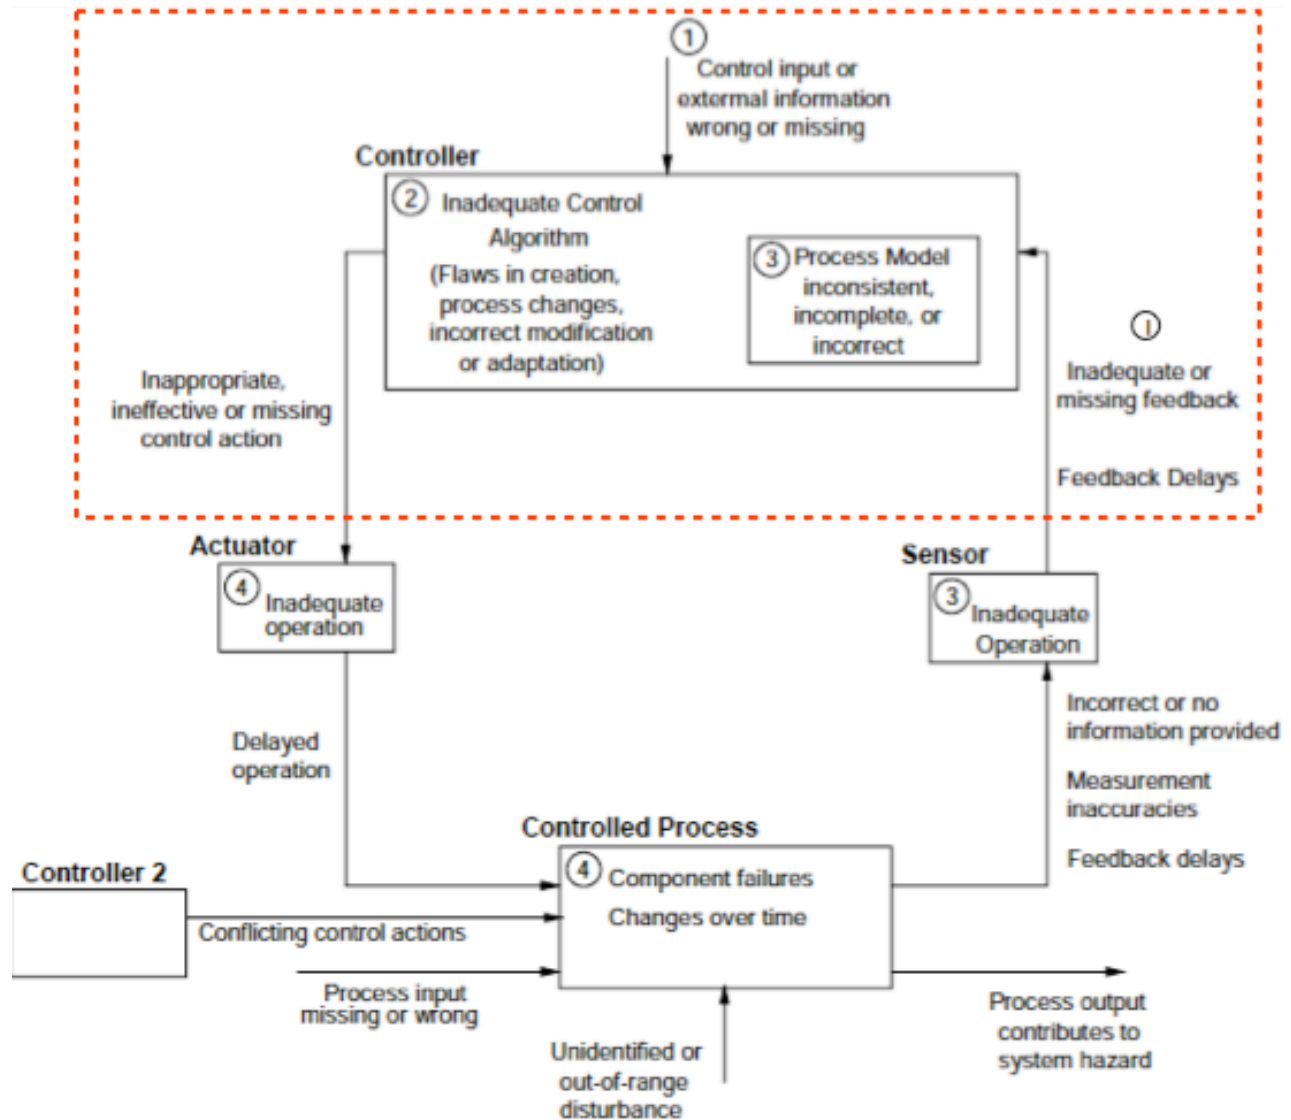
Figure 36. Original Human-Controller Hazard Analysis [Leveson 2012]

As Figure 36 shows, these four categories revolve around inadequate feedback, wrong control input or external information, inconsistent process models, and an inadequate control algorithm.

129

Applying the updated human-controller model from the previous section, the causal-factor analysis now explicitly related to a *human controller* is seen in Figure 37.

As a whole, the updated human-controller analysis in Figure 37 includes five distinct causal-factor categories that largely revolve around the process states for any given unsafe control action. The five categories above directly stem from Figure 35 through the distinct stages of "feedback," human cognition ("detection and interpretation," "mental models," and "decision-making"), and the "affordance" of action.

The first update in Figure 37 resides in the analysis of flawed feedback as denoted by the (1), or first causal-factor category. This updated model combines all feedback from Figure 35 (i.e., from the displays, sensors, and sensory perception) and *centers this causal-factor category of feedback on the process states themselves*. It is important to note that this first category analyzes the feedback in and of itself *before* it reaches the human controller and indicates whether or not that feedback, which is used to update the human controller about the process states for a given unsafe control action, is conflicting, incorrect, missing, delayed, or unrefreshed. As the next sub-section details, the analysis of this first category will methodically address each component of feedback by focusing on each high- and lower-level process state. Furthermore, this first category also includes causal factors related to feedback from the affordance and implementation of an unsafe control action despite there being no direct link between affordance, action, and feedback. After this feedback reaches the human, there are four more causal-factor categories that could lead to unsafe control.
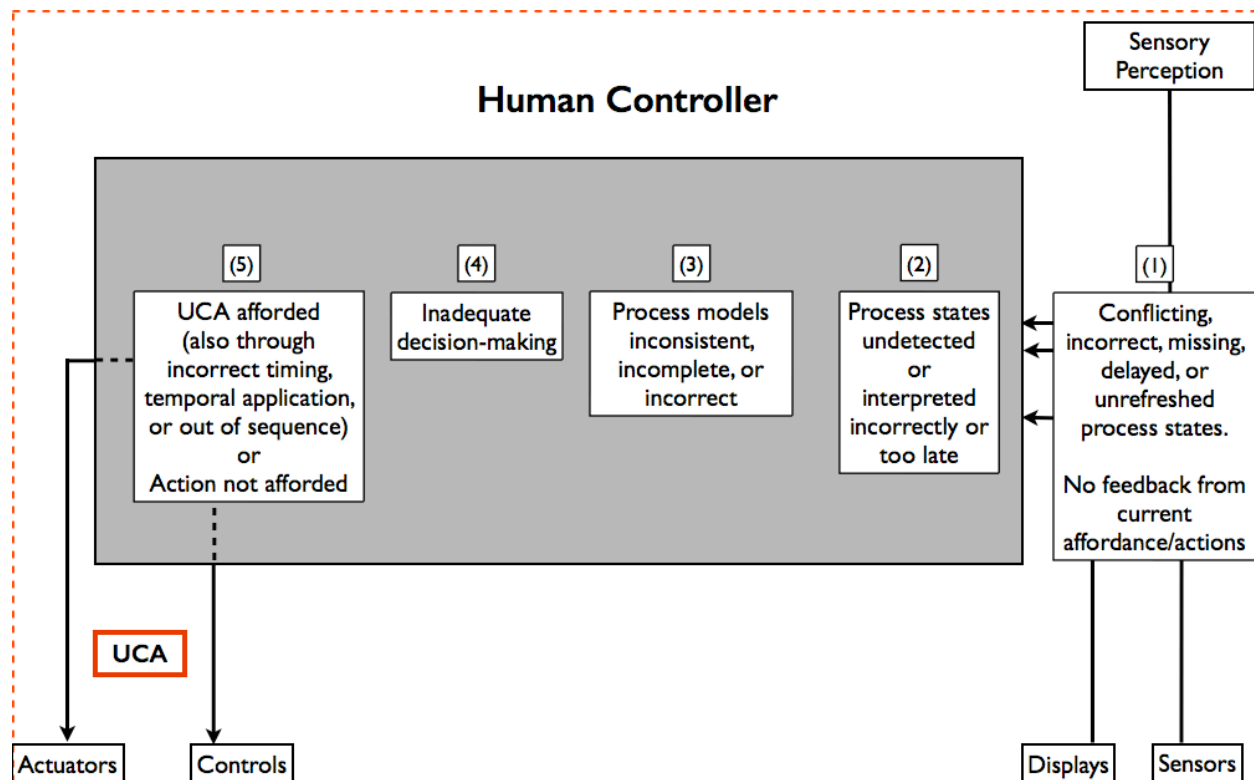

Figure 37. The Updated Human-Controller Analysis

The second category of causal factors seen in Figure 37 focuses on the flawed detection and interpretation of feedback (2) *after* it reaches the human. In contrast to dissecting the feedback itself, this new stage explicitly analyzes potential flaws in the human detection and interpretation of feedback that could lead to unsafe control. In this manner, the distinction between feedback and the human detection and interpretation of that feedback is clarified to better organize the ensuing causal factors. Thus, in order to achieve system safety, the human controller must accurately detect and interpret the process states as they exist in reality. Anything other than the accurate detection and interpretation of the process states through feedback (displays, sensors, and the human's own sensory perception) will, when paired with a worst-case set of environmental conditions, lead to an accident.

While the process states are useful in determining *what* feedback to analyze, another reason for including the detection and interpretation category (2) is to address *how* the human receives the feedback. *What* feedback is flawed (1) is just as important as *how* the detection and interpretation of that feedback can become flawed (2). While the design of the system around the human controller is tightly coupled with system safety (e.g. the presentation of feedback, the layout of controls), this approach to analyzing *how* must be treated with caution because the goal of STPA is hazard analysis and not design. This new category will, however, accomplish hazard analysis by better informing the designers of how the feedback can become undetected or misinterpreted through the use of process states as the next section will elaborate.

The next two categories of inconsistent process models (3) and flawed decision-making (4) largely remain unchanged from the previous human-controller analysis in Figure 36. Aside from the updated semantics in (4), the *method* by which flawed process models (3) are generated has changed. In the previous human-controller analysis, there was no explicit means of generating the causal factors in (3), but in this updated analysis, all causal factors in (3) will derive directly from the process states—detailed in the next sub-section.

Following the same logic as Figure 35, the last causal-factor category is that of inappropriate affordance (5). In this category, the causal factors that can lead to an accident stem from the inappropriate affordance or inaffordance of action that leads to an unsafe control action. There is also the unseen link between affordance and feedback that might become severed and contribute to an unsafe control action. In either case, this updated human-controller analysis in now specifically includes a new category that magnifies these issues in affordance.

Altogether, the biggest changes in appearance in Figure 37 were explicitly identifying the controller as human, grouping the feedback into one coherent category, and adding the "detection and interpretation" and "affordance" categories. More importantly, much of the subsequent causal-factor analysis will now be framed around the system and environmental properties that are necessary for system safety—the process states. Despite these improvements in appearance, the strongest and most distinctive update to this model is not in the new categories themselves, but in the method by which the causal factors are generated.

9.3   *An Applied Example using In-Trail Procedure (ITP)*

In 2008, a safety analysis report titled *DO-312 (Safety and Interoperability Requirements Document for the In-Trail Procedure in Oceanic Airspace (ATSA-ITP) Application)* was

completed by the Radio Technical Commission for Aeronautics (RTCA) to provide "the minimum operational, safety, performance, and interoperability requirements" for ITP [RTCA, 2008]. It was in response to this original safety analysis on ATSA-ITP (Airborne Traffic Situational Awareness In-Trail Procedure)—an analysis based on fault trees and probabilistic assessments— that a 2012 STPA report (*Safety Assurance in NextGen*) was conducted to compare traditional safety methods with one based on the systems-theoretic approach of STPA [Fleming et al 2012]. This section will detail the approaches to the human controller in both DO-312 and the 2012 STPA report before applying the updated methodology to further extend the STPA analysis.

As part of the FAA's upcoming NextGen architecture, the primary objective of In-Trail Procedure is to enable aircraft to achieve flight level changes on a more frequent basis to improve flight efficiency. This new procedure reduces separation minima through enhanced aircraft position and velocity reporting by means of ADS-B—a GPS-enabled aircraft transponder—which thereby allows aircraft to perform climb-through or descend-through maneuvers in procedural airspace where no radar services are provided. Onboard the flight deck is also a new set of avionics specifically intended to assist the flight crew in conducing In- Trail Procedure, referred to as the ITP Equipment. ITP itself is comprised of six different maneuvers [RTCA, 2008]:

1. A following climb
2. A following descent
3. A leading climb
4. A leading descent
5. A combined leading-following climb
6. A combined leading-following descent

An example of an ITP following climb is seen in Figure 38.
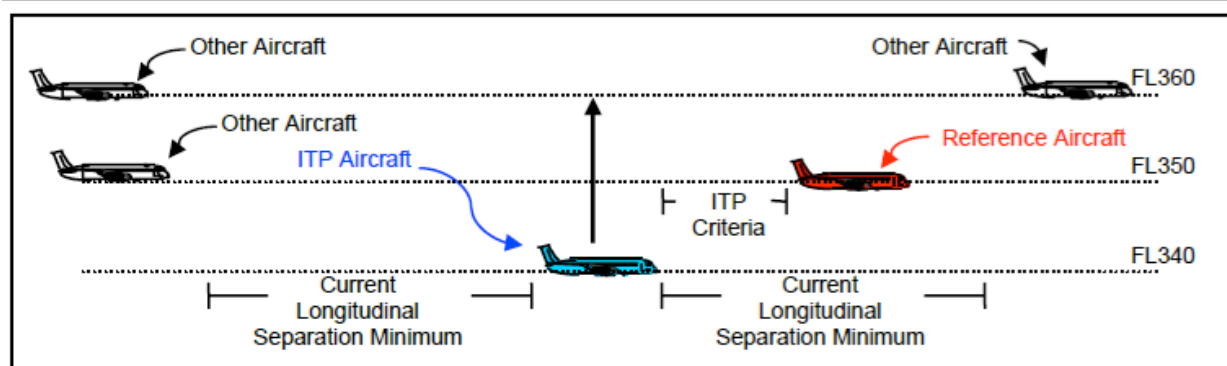
Figure 38. ITP Following Climb [RTCA, 2008]

In this example, both the ITP aircraft and the reference aircraft are equipped with ADS-B receive and transmit functions (ADS-B In and Out), so both aircraft as well as Air Traffic Control (ATC) are aware of the position and velocity data of the aircraft. With the known position and velocity of the reference aircraft through ADS-B, the flight crew of the ITP aircraft in this example is then responsible for four main items before the maneuver can be executed:

1. Checking and validating a set of ITP criteria for the following-climb maneuver with the assistance of the ITP Equipment.

2. Checking that the airspace is clear of conflicts with traffic, poor weather, turbulence, etc.

3. Requesting and receiving clearance from ATC to execute the following-climb maneuver.

4. Verifying the ITP criteria for the following climb maneuver with the assistance of the ITP Equipment.

After these four items have been met, it is then appropriate for the flight crew to execute the following-climb ITP maneuver to the requested flight level. Since this maneuver takes place in procedural airspace, the flight crew plays the largest role in ensuring the maneuver is executed in a safe state—with ATC serving as a balance in verifying the ITP criteria and maneuver for the flight crew. It was due to the heavy involvement of the flight crew and hence the human controller in the execution of this ITP procedure that this example was ripe for human-controller analysis.

9.4 Human-Controller Analysis using STAMP Accident Model

Shifting now to STPA, this section will focus on the human-controller analysis in the 2012 STPA report and then apply the updated human-controller analysis from Figure 37 to further extend the results. Before diving into the STPA Step 2 analysis of the human controller, however, it is imperative to define the basic accidents, hazards, and safety control structure from the System Engineering Foundation as well describe the specific unsafe control actions that will be analyzed. As detailed elsewhere, the accidents under consideration are human death or injury and the high-level hazards are [Fleming et 2012]:

H-1: A pair of controlled aircraft violate minimum separation standards.

H-2: Aircraft enters unsafe atmospheric region.

H-3: Aircraft enters uncontrolled state.

H-4: Aircraft enters unsafe attitude (excessive turbulence or pitch/roll/yaw that causes passenger injury but not necessarily aircraft loss).

H-5: Aircraft enters a prohibited area.

9.4.1 Process States and Unsafe Control Actions

The unsafe control actions chosen for this analysis all derived from the flight crew's execution of In-Trail Procedure. Due to a somewhat ambiguous definition in the RTCA description, however, "execute ITP" was refined to include the stages of the "initiation" and "continuation of ITP," and omitted the "termination of ITP" in this example. Following the systematic method of STPA, the process states (**PS's**) for the initiation and continuation of ITP had to first be defined:

PS 1: ITP criteria

PS 2: ATC clearance

PS 3: Airspace model

The first Process State of ITP criteria are actually "[a set of 10] conditions that must be satisfied prior to initiating or executing an ITP clearance," and include criterion such as the ITP distance between aircraft, vertical speed requirements, Mach differentials, data integrity, etc. [RTCA 2008]. As ITP is currently designed, these ITP criteria must be checked and validated by the flight crew before requesting ATC clearance and then validated again prior to executing the ITP maneuver. The second Process State of ATC clearance is a system condition where ATC clearance is either approved or not. Since this example centers on the flight crew, the ATC clearance in this scenario is assumed to be correct—regardless of being approved or not. Finally, the third high-level Process State of the airspace model is that state of the airspace outside of the immediate ITP maneuver. For ITP to be executed in a safe state, the ITP aircraft of Figure 14 must not violate minimum separation standards with other aircraft as well as not fly into an unsafe atmospheric region and become uncontrollable—a direct relation to high-level hazards H.1-H.4. Given the 10 criteria within Process State 1 and the traffic/weather conditions within PS 3, both can be broken down into lower-level process states as shown in Table 28.

Table 28. Hierarchy of Process States

High-level Proc. State	1. ITP Criteria (met or not)	2. ATC Clearance (approved or not)	3. Airspace Model (clear or not)
Lower-level Proc State	**1.1 Climb/Descent rate** (Y/N) **1.2 ITP distance*** (Y/N) **1.3 Ground speed differential*** (Y/N) **1.4 Mach differential** (Y/N) **1.5 Reference a/c maneuvering or expected to** (Y/N) **1.6 Vertical distance reqs** (Y/N) **1.7 Ownship data integrity** (Y/N) **1.8 Reference a/c data integrity** (Y/N) **1.9 Same track criteria*** (Y/N) **1.10 Requested flight level**	None	**3.1 Weather clear for ITP** (Y/N) **3.2 Clear of other traffic** (Y/N)

	correct (Y/N)		

The first important detail in Table 28 is that the high-level process states have two fundamental values: valid or invalid for ITP. It follows, then, that In-Trail Procedure is only valid when the ITP criteria have been met, ATC clearance has been approved, and the airspace is clear— and if any of these states are invalid ITP cannot be executed in a safe state. Secondly, each high- level process state can only be valid if every single lower-level process state within is valid (i.e., the ITP criteria (Process State 1) can only be valid if Process State 1.1-1.10 are all valid), which the "(Y/N)" in Table 9 is used to indicate. Lastly, the asterisks (*) next to Process State 1.2, 1.3, and 1.9 show which lower-level process states the ITP Equipment is planned to calculate on its own [RTCA 2008].

From this process-state hierarchy, the unsafe control actions selected for further causal analysis were then identified and listed as shown in Table 29.

Table 29. Selected Unsafe Actions for Further Causal Analysis

Control Action	Unsafe Control Action with Process States
Initiate ITP	ITP initiated when **ITP criteria** (**PS 1**) have not been met
	ITP initiated when **ATC clearance** (**PS 2**) has not been approved
	ITP initiated when **airspace model** (**PS 3**) is not clear
Continue ITP	ITP continued with inappropriate **ITP criteria** (**PS 1**)
	ITP continued with revoked **ATC clearance** (**PS 2**)
	ITP continued with an **airspace model** (**PS 3**) that no longer permits ITP

All of the UCA's in Table 29 were framed around the high-level process states and used for parallel causal (hazard) analysis. That is, the human-controller causal factors the hazard analysis were generated using all six of these unsafe control actions at once and were not analyzed independently.

9.4.2   Human-Controller Analysis found Using Existing STPA

From the UCA's described in Table 30, the 2012 STPA report then generated the causal factors related to the human controller through the original template of Figure 36. All causal factors that were found to relate to the human controller are tabulated in Table 30.

Table 30. Original Causal Factors Related to the Human Controller [Fleming et al 2012]

Hazard	H.1-H.4

| Unsafe Control Action | - **ITP initiated** when any of **PS 1-3** are not met, approved, or clear for ITP |
| | - **ITP continued** when any of **PS 1-3** are no longer met, approved, or clear for ITP |

Process Model Link	Extended Causal Factors
Inadequate control algorithm	Flight Crew: - does not correctly check that ITP is appropriate (that normal FL change could occur) - does not check that all ITP criteria are met - begins executing ITP prior to receiving approval - delays in executing ITP after receiving approval - does not re-verify that conditions have changed from when they were originally checked after receiving approval - does not confirm the established flight level after finishing the maneuver
Process model inconsistent	Flight Crew believes: - that their climb/descent capability is greater than it is - it has all ADS-B data for local traffic - ADS-B data to be accurate when it is not - ITP criteria (speed, distance, relative altitude, relative angle) to be different than it is - communication protocols with ATC to be different than they are - communication with nearby aircraft to be different than they are - in a different understanding of individual crew member responsibilities - weather/turbulence to be better than it is - ITP request to be approved when it is not - ATC approval to be recent when it is old
Inadequate Sensor Operation	- Flight crew does not understand or correctly apply the ITP data from the ITP equipment
Control input or external information wrong or missing	- Flight crew lacking information from ATC - ATC approval not on communication channel that FC is monitoring - ITP Equipment provides criteria data too late - ITP Equipment gives incorrect or ambiguous state information
Inadequate or missing feedback	- Change in the velocity/altitude/bearing of ownship not displayed to pilot - Change in the velocity/altitude/bearing of nearby ship not displayed to pilot

	- Proper aircraft identifier of nearby aircraft not displayed to pilot - FC does not receive communication from ATC - Flight crew does not receive local traffic information from ADS-B

As with any causal-factor generation in STPA, the Step 2 analysis must first connect to a specific set of unsafe control actions (or single unsafe action) as well as the relevant high-level hazards, both of which are detailed in the first two rows of Table 29. The hazards all trace back to H.1-H.4, and the UCA's fundamentally translate to "the initiation or continuation of ITP when any of the high-level process states are invalid for ITP." Located underneath this is the summary of all categories or "process model links" of the causal-factor analysis that trace back **to** Figure 36**.**

In this report, the causal factors in Table 30 that were related to the flight crew stem from the categories of: inadequate control algorithms, process model inconsistencies, inadequate sensor operation, wrong external information or control inputs, and inadequate or missing feedback. To reiterate, the identified causal factors are not a comprehensive list of all those that were generated in the 2012 STPA report, but rather all those causal factors that specifically involved the flight crew for these UCA's. Notice that there is no relation to estimating human reliability or linear combinations of human failure that result in a hazardous state. The basic control loop template of Figure 36 instead views the human controller (the flight crew) *within* the context of the system and controlled process occurring in the ecology to generate causal factors that could lead to accident scenarios. Furthermore, these causal factors lead to safety constraints (or requirements) that can be handed to system designers to ultimately design a safer system—one that does not operate in a hazardous state—and offer much more than random human error probabilities can convey. The methodology for using Figure 36 in generating the causal factors in Table 30, however, is not clearly defined in STPA.

9.4.3   Extending the Human-Controller Analysis

In pursuit of extending this analysis, the new human-controller methodology was applied to the unsafe actions from Table 29 to generate an entirely new set of causal factors. Following the process outlined in the previous sub-section, this extended set of causal factors were derived according to the five categories from Figure 37. The causal-factor analysis first started with the inadequate process models (3), shown in Table 31.

Table 31. Inconsistent Process Model (3) Categories

Hazard	H.1-H.4
Unsafe Control Action	- **ITP initiated** when any of **PS 1-3** are not met, approved, or clear for ITP - **ITP continued** when any of **PS 1-3** are no longer met, approved, or clear for ITP

Process Model Link	Extended Causal Factors
(3) Process models inconsistent, incomplete, or incorrect	Flight Crew believes: - **ITP Criteria** (**PS 1**) has been met when it has not - **ATC clearance** (**PS 2**) to be valid when it is not - **airspace model** (**PS 3**) to be clear when it is not

In this category, any of the high-level process states are believed by the flight crew to not be in a hazardous state when in fact, they are. The biggest change in this category of flawed process models (3) stems from centering the causal factors on the abstract high-level process states, as this structure pinpoints exactly what beliefs on behalf of the flight crew may lead to an accident—an approach that would work equally as well for an automated controller. In addition, the three incorrect beliefs about the high-level process states are all part of a logical "OR" statement, so an accident can occur if the flight crew believes any one of these three things—an idea of logical simplification that may help future analyses.

The next category of inadequate decision-making (4) reflected no changes from the 2012 STPA report, aside from the change in title, so the next category of analysis shifted to the inappropriate affordance of ITP (5) as seen in Table 32.

Table 32. Inappropriate Affordance of ITP (5) Causal Factors

Hazard	H.1-H.4
Unsafe Control Action	- **ITP initiated** when any of **PS 1-3** are not met, approved, or clear for ITP - **ITP continued** when any of **PS 1-3** are no longer met, approved, or clear for ITP
Process Model Link	**Extended Causal Factors**
(5) ITP inappropriately afforded	- Flight Crew inappropriately affords the initiation of ITP or continues to afford ITP, through a slip or mistake, and isn't made aware of this through feedback

This category of inappropriate affordance (5) in **Table 13** captured those ways by which the unsafe **initiation** and **continuation of ITP** could occur and to how the flight crew could not be made aware of this through feedback. This category also highlighted a missing link between affordance, action, and feedback—an unseen link that may bypass the controlled process and/or automation altogether and instead feed directly back to the human.

138

Following this, the causal-factor analysis then shifted to the treatment of flawed feedback (1) *before* it reaches the flight crew through the structure of the process states. As Table 33 shows, each high-level process state and its lower-level hierarchy were methodically addressed.

Table 33. Flawed Feedback (1) Causal Factors

Hazard	H.1-H.4
Unsafe Control Action	- **ITP initiated** when any of **PS 1-3** are not met, approved, or clear for ITP - **ITP continued** when any of **PS 1-3** are no longer met, approved, or clear for ITP
Process Model Link	**Extended Causal Factors**
(1) Conflicting, incorrect, missing, delayed, or unrefreshed process states. *No traceability to current affordance/actions*	Any of the **ITP criteria (PS 1.1-1.10)**: - are incorrect or missing - aren't refreshed in the appropriate amount of time - are in conflict which leads to an ambiguous **ITP criteria (PS 1)** **ATC clearance (PS 2)**: - is incorrect or missing - isn't provided in the appropriate amount of time - no longer remains valid (i.e. not refreshed in the appropriate amount of time) Either **Airspace model variable (PS 3.1 or 3.2)**: - is incorrect or missing - isn't refreshed in the appropriate amount of time - is in conflict which leads to an ambiguous **Airspace model (PS 3)** - There is a conflict between **ITP criteria**, **ATC approval**, and the **airspace model** (i.e. a conflict between **PS 1**, **PS 2**, and **PS 3**)

Using the ITP criteria (Process States 1.1-1.10) from Table 33 as an example, any of the feedback used to represent the lower-level process states (PS 1.1-1.10) can be, incorrect, missing, or unrefreshed in the appropriate amount of time. Not only this, but if any of the lower-level ITP criteria (Process States 1.1-1.10) conflict with themselves through feedback (e.g., ITP distance readings from the Traffic Collision Avoidance System (TCAS) and the ITP Equipment conflict

with each other and lead to a conflicted ITP distance (Process States 1.2)), the overall state of the high-level ITP criteria (Process State 1) will be ambiguous, unknown, or undefined—i.e., conflicted. This same approach was then applied to the other process states of ATC clearance (Process State 2) and the airspace model (Process States 3.1-3.2). Another source of causal factors in flawed feedback (1) comes from conflicts between the high-level process states themselves which could lead to flawed mental models and decision-making (e.g., ATC approves ITP clearance even though the ITP criteria are not met and the airspace model is not clear for ITP. This could be further refined by asking: how could the flight crew initiate or continue ITP based on ATC's clearance (Process State 2) when it conflicts with Process State 1 and Process State 3?). A notable discovery in this process is that *conflicts* in feedback are now easily identifiable. Moreover, another important distinction in this approach is that *all* feedback to the flight crew is now categorized with respect to the process states, an approach that could also equally apply to an automated controller. In this manner, any number of ways by which the high- and lower-level process states become incorrect, missing, unrefreshed, or conflicted *through feedback* can be assessed—a large advantage for a system that has not even been fully designed yet.

Finally, the causal factors in the last remaining category of flawed detection and interpretation (2) were generated and are shown in Table 34.

Table 34. Flawed Detection and Interpretation (2) Causal Factors

Hazard	10. H.1-H.4
Unsafe Control Action	- **ITP initiated** when any of **PS 1-3** are not met, approved, or clear for ITP - **ITP continued** when any of **PS 1-3** are no longer met, approved, or clear for ITP
Process Model Link	**Extended Causal Factors**
(2) Process states undetected or interpreted incorrectly or too late	Any of **ITP criteria** (**PS 1.1 - PS 1.10**) OR their changes/updates: - are not detected - are not interpreted correctly and leads to inaccurate or conflicting understanding of the **ITP criteria (PS 1)** - take too long to detect and interpret correctly - require too much attentional demand to detect and interpret correctly **ATC clearance (PS 2)** or any change or update: - Anything but ATC clearance is detected and interpreted as a clearance - A revoke of ATC clearance is not detected and interpreted correctly Either **Airspace variable (PS 3.1 or 3.2)**: - is not detected - is not interpreted correctly and leads to

	inaccurate or conflicting understanding of the **Airspace (PS 3)**
	- takes too long to detect and interpret correctly - requires too much attentional demand to detect and interpret correctly
	- There is a conflict between the flight crew's interpretation of the **ITP criteria**, **ATC approval**, and the **airspace model** (i.e. a conflict between **PS 1**, **PS 2**, and **PS 3**)

The formation of these causal factors followed the same pattern as the last category of flawed feedback (1) by first starting with each high-level process state and its associated hierarchy. For example, any of the lower-level ITP criteria (Process States 1.1-1.10) OR their changes and updates may not be detected and interpreted correctly in the appropriate amount of time, they may take too much attentional demand to do so, or they may be interpreted incorrectly which leads to an inaccurate or conflicted understanding of the high-level **ITP criteria** (**Process State 1**) itself. An important understanding here is that while the feedback may be correct before it is detected by the flight crew, it matters none if the they cannot correctly interpret the feedback and understand the process states as they exist in reality. This misunderstanding of feedback is a causal factor that can easily lead to flaws in the process models (3) of the flight crew and an eventual unsafe initiation or continuation of ITP when any of the high-level process states are invalid. Furthermore, just as the high-level process states could conflict with each other in flawed feedback (1), so too can the flight crew's interpretation of those high-level process states as Table 34 lists.

10.1.1 Discussion of the Extension

From the original causal factors in the 2012 STPA report, the extension in this new approach changed the underlying structure and methods by which the causal factors were generated, which resulted in a longer and more structured list of causal factors and scenarios related to the flight crew. The backbone of this change was through the integration of the process states in the first three causal-factor categories of flawed feedback (1), flawed detection and interpretation (2) of feedback, and inconsistent process models (3). Much like the meaningful "state" variables in ecological situation are at the heart of what matters for effective human control, so too are the process states in STPA at the heart of what matters for safe control. Leveraging this principle, the entirety of feedback that reaches the human operator, to include their own sensory perception, is now analyzed in one specific category. An interesting finding is that the method for analyzing the categories of flawed feedback (1) and inconsistent process models (3) through the process states applies equally as well to an automated controller. Although the original intention was to improve the human-controller analysis, this is application to the automated controller is a welcome finding.

In addition to leveraging process states to frame parts of the causal-factor analysis, this new approach grounded in Figure 37 revealed causal factors that had previously been difficult to identify—those related to *conflicts* in feedback. In the original methodology there were no clear means of establishing conflicts between feedback, or more importantly, *which conflicts mattered* (with respect to the process states). If the feedback or the flight crew's interpretation of it leads to a conflicted model of the airspace, or if the ITP distance is ambiguous, for example, the system can be designed to warn them or prescribe a course of action if this does occur. It is the entire

purpose of STPA hazard analysis to hand the designers safety constraints that come from detailed causal factors just like this, but in light of all five of these categorical analyses it is necessary to discuss the ramifications of this updated analysis.

10.1.2 Potential Future Work on this Extension

Despite the improvements in clarity and organization in this updated analysis, there still remain some limitations to the new approach. The first criticism of this new analysis is that some of the generated causal factors can still be vague. For example, the category of affordance (5) highlighted a case where the flight crew inappropriately afforded the initiation or continuation of ITP and was not made aware of it through feedback. Aside from listing this as a causal factor, there is no further explanation to this statement. Although this is certainly an accurate criticism, it is important to remember that this analysis was conducted on a system *that has not even been designed yet*, so some of the causal factors are expected to be vague. For In-Trail Procedure, the automation has not been fully developed and the aircraft used in conducting ITP will vary—thus the affordance of ITP in each aircraft has still yet to be defined. With this analysis, however, the safety constraints generated from these updated causal factors will help guide designers to construct a safe system in order that this causal factor is prevented.

Another criticism of this updated analysis is the poor link between affordance, action, and feedback. Aside from merely stating "no link between affordance, action, and feedback" as a type of causal factor, this new approach offers no more guidance. An assumption in STAMP theory is that all control actions go through the controlled process or automation before it eventually reconnects to the human controller through feedback from the sensors and displays. In most human-controlled systems this is not always the case, and the human operator may receive direct feedback from their affordance or action, such as when a pilot physically moves the yoke, for example. Perhaps there is another feedback arrow that comes directly the human controller's action, a possibility that should be explored in the future. Perhaps most importantly, the extended analysis at identifies this as a causal factor, whereas other approaches neglect such factors.

11. Relationship of STAMP, STPA, and STECA to Safety Management Systems

The requirements of a Safety Management System (SMS) comprise multiple layers of requirements originating with international policy set out by ICAO in agreement with member States, through States national policy (FAA, EASA etc.) to aviation service provider organizations. The ICAO and States guidance is not prescriptive. It sets out a series of objectives (i.e. what is required) without specifying the means by which that is achieved. This is essential since it is not feasible for States to adequately specify an SMS for each organization due the wide variety of sizes, each with unique functions, structure and operations. As a result, the implemented SMS is intended to be developed by the organization, for its own use and according to its specific needs. Each organization will seek approval of its SMS from the appropriate national regulator. An SMS consists of all the organization specific safety assurance policies and procedures that implement the high level international and States requirements along with a supporting quality assurance process to assure that the documented processes are applied consistently. This approach mirrors the ISO-9001 approach of "say what you, do what you say" that has been widely applied in quality management systems for many years.

11.1 Safety Management System Overview

At the international policy level, ICAO has developed Annex 19 to the Convention on International Civil Aviation [ICAO, 2013a] and a Safety Management Manual [ICAO, 2013b]. These documents provide the high level guidelines for States to implement the objectives and requirements of the ICAO Global Aviation Safety Plan (GASP) [ICAO, 2013c].

The FAA defines an SMS as "the formal, top-down business-like approach to managing safety risk, which includes a systemic approach to managing safety, including the necessary organizational structures, accountabilities, policies and procedures". In general States require that aviation service providers document their safety policies and procedures in an SMS Plan and submit it to the national regulator for approval.

The FAA has an initiative to require aviation service providers to create and get approval for a SMS. The FAA convened an Aviation Rulemaking Committee (ARC) to provide recommendations [FAA, 2010a] and the resulting framework is outlined in [FAA, 2010b]. The FAA intends that SMSs be provided by all classes of aviation service providers as listed below. FAA has issued AC 120-92B [FAA, 2015] to provide guidance on the required content of an SMS and internal Orders [FAA, 2013a] [FAA, 1998]. Aviation service providers are classified into four broad groups.

> (1) Air Operators, Maintenance and Repair Organizations, Flight Training organization
> (2) Airports
> (3) Air Traffic
> (4) Design and Manufacturing (D&M)

Following the usual regulatory process, the FAA works with the aviation industry through an Aviation Rulemaking Advisory Committee (ARAC) to develop the text of a Proposed Rule that

is approved and issued by an ARC in a Notice of Proposed Rulemaking (NPRM) published in the Federal Register. After a public comment period, the proposed Rule becomes final. The FAA will then publish guidance material in the form of Advisory Circulars and develops Orders to inform FAA staff on how to assess an applicant's conformance to the Rule.

At this time, a Final Rule in place only for Part 121 Air Operators (Group 1) {Fed. Regist., 2015]. The FAA has issued an NPRM for Group 2 (airports) {Fed. Regist., 2010] and an associated advisory circular [FAA, 2007]. For Group 3 (Air Traffic), the FAA has developed an SMS manual [FAA, 2014a] and internal Orders [FAA, 2012a] [FAA, 2013b]. The FAA acts as the approval authority for organizations operating in the US. In the case of Air Traffic, the FAA Air Traffic Oversight Service acts as the approver.

At the present time there is no NPRM published in the Federal Register for design and manufacturing operations (Group 4, D&M). D&M includes aircraft and equipment manufacturers. An ARC for design and manufacturing operations completed their activities and submitted a report to the FAA in October 2014 [FAA, 2014b]. A pilot project has been completed [FAA, 2011a] reported in [FAA, 2012b] and a guidance document [FAA, 2011b] produced.

11.2   Structure of an SMS

An SMS consists of four basic components which are common to the SMS of all types of aviation service providers. The structure below is taken from [FAA, 2011b] but mirrors ICAO SARPs Annexes 6 [ICAOd] and 19 [ICAO, 2013a] and [FAA, 2015].

Safety Policy and Objectives (Component 1.0)

Effective management systems must define policies, procedures, and organizational structures to accomplish their goals. The SMS Framework's Safety Policy and Objectives Component outlines expectations in the Elements below, which in turn provide the foundation for the functional SMS Components 2.0 and 3.0 (Safety Risk Management and Safety Assurance).

- Safety Policy (Element 1.1)
- Management Commitment and Safety Accountabilities (Element 1.2)
- Key Safety Personnel (Element 1.3)
- Emergency Preparedness and Response (Element 1.4)
- SMS Documents and Records (Element 1.5).

Safety Risk Management (Component 2.0)

Safety Risk Management (SRM) is a formal system of hazard identification and analysis and risk control (sometimes termed "mitigations") used to assess systems at both the organizational and product levels. SRM Framework Elements are essential in controlling risk to acceptable levels and their subordinate Sub-Elements are:

- Hazard Identification and System Analysis (Element 2.1)
- System description and analysis (Sub-Element 2.1.1)
- Identify hazards (Sub-Element 2.1.2)
- Risk Assessment and Control (Element 2.2)
- Analyze safety risk (Sub-Element 2.2.1)

- Assess safety risk (Sub-Element 2.2.2)
- Control/mitigate safety risk (Sub-Element 2.2.3).

Safety Assurance (Component 3.0)

Once SRM controls are identified and employed, an organization must ensure that the SRM-designed and implemented controls continue to be implemented as intended and are effective as the environment changes. The Safety Assurance (SA) function provides for this, using system safety and quality management concepts and sub-elements. SA Framework Elements for assuring safety and the subordinate Sub-Elements are:

- Safety Performance Monitoring and Measurement (Element 3.1)
- Continuous monitoring (Sub-Element 3.1.1)
- Internal audit (Sub-Element 3.1.2)
- Internal evaluation (Sub-Element 3.1.3)
- Investigation (Sub-Element 3.1.4)
- Employee reporting and feedback system (Sub-Element 3.1.5)
- Analysis of data (Sub-Element 3.1.6)
- System assessment (Sub-Element 3.1.7)
- Management review (Sub-Element 3.1.8)
- Management of Change (Element 3.2).

Safety Promotion (Component 4.0)

The organization's upper management must promote safety as a core value with practices that support a sound safety culture. The organization must make every effort to communicate its goals and objectives as well as the current status of the organization's activities and significant events. The Safety Promotion Component provides the SMS Framework expectations for establishing and implementing these functions through the following Elements:

- Competencies and Training (Element 4.1)
- Personnel expectations (competence) (Sub-Element 4.1.1)
- Training (Sub-Element 4.1.2)
- Communications and Awareness (Element 4.2).

11.3 Potential Role of STPA in SMS

An SMS is not a safety assessment process; it is high level process and structure within which a variety of safety practices (including hazard and safety assessment methods such as STPA and STECA) can reside. STPA fits into the Safety Risk Management (SRM) pillar of the SMS Manual v4.0 [FAA, 2014] under the steps requiring the identification of hazards and safety risk mitigation. STPA also fits in by providing a way to describe the system in the system analysis step of the SRM process. The STPA functional control diagram includes all the things currently described in this SMS requirement.

The more complete description provided in the STPA functional control diagram will be needed for system design and maintenance and particularly for any future changes that are made. In the extensive use of STPA in other industries, one of the most common responses from users is that the functional control diagram produced for STPA is the best and most useful documentation

they have about the functional design of the system, even in a decades long development program near its end, as was the case for the application of STPA on the U.S. Missile Defense System. Most industry documentation, particularly graphical documentation, focuses on the physical system design and not the functional design.

A third place where STPA would be extremely useful is in the Safety Assurance pillar. The STPA causal analysis can be used to identify what needs to be assured, including audit procedures to verify that safety is not degrading as behavior changes over time.

STAMP does have a slightly different definition of hazard than included in the SMS Manual. The official definition in the SMS manual is "a condition that is a prerequisite to an accident or incident" The drawback of that definition is that there are a very large number of conditions that precede an accident. Aircraft being in the airspace is prerequisite to an accident or incident, but we cannot eliminate that condition, i.e., not allow any planes in the airspace. The definition used in STPA restricts hazards to be hazardous conditions or states, such as a violation of separation standards. These conditions, once they are identified, can be eliminated or controlled in the system design and operations. All prerequisites to an accident (the SMS definition) cannot be considered (and do not need to be) as they include almost all conditions that occur during normal operations.

In practice, we suspect that the actual hazards identified in any SMS hazard identification process will be those that fit the STPA definition. Otherwise, such a listing would be impossible. This was the case for an STPA safety analysis of the In-Trail Procedure (ITP) [Fleming et al, 2012]. In fact, official ITP analysis identified fewer hazards than STPA because it only considered component failures. The same is true for the techniques currently listed in Appendix G of the ATO SMS Manual. As such, the hazards identified by the Appendix G methods will be a proper subset of those identified by STPA (assuming both are competently executed).

There may be implications for the use of STPA on the process of assessing risk in SMS. Because human errors and non-failure scenarios can be identified by STPA, it will not be possible in many, or maybe even most, cases to assess a probability for these hazards. A different type of assessment will be needed. One solution to this dilemma is simply to ignore the scenarios for which probabilities cannot be derived, which is effectively what is being done now because those scenarios are never identified even though they exist. The result is (and would be) an inaccurate value. Unfortunately, there is no way to evaluate the accuracy of such probabilistic risk assessments (without waiting a thousand years to see what happens) so they are usually just accepted as true without validation. Non-probabilistic assessment methods will be needed, but this conclusion has little to do with STPA and everything to do with the nature of hazards in complex, tightly coupled, distributed systems. The possibility of using qualitative methods for risk assessment is included in the SMS Manual so the change necessary will be in practices used rather than in the definition of and requirements for risk assessment in SMS.

STPA has no impact on the SMS processes involving identifying mitigations and controls. STPA is an analysis technique, not a design technique. If used on a design that already exists, STPA does not tell the engineers how to redesign the system but it does provide important information including, whether adequate controls already exist or if new ones are needed. The method can also potentially be applied to various alternatives to see which one(s) are preferable from a safety standpoint. If used early in the design process, as in this report, STPA can generate the safety design requirements for the system being considered so that better decisions can be made during

concept formation and early design. Assessing safety after the design is complete usually limits the types of hazard mitigations that are possible and results in extremely costly (both in terms of money and time) rework. Designing safety in as the design concepts and decisions are made can result in safety costing very little or even nothing extra.

Another potential use for STPA in the SMS process is to evaluate the independence and safety impact of potential changes and upgrades in ATC capabilities. The Preliminary Safety Analysis defined in the SMS Manual requires determining whether a change can impact safety and therefore requires more extensive analysis. This process is actually much more difficult than people think. An informal impact analysis (and even one using the traditional hazard analysis techniques) can only identify the direct relationships between the changes proposed and the current system. Indirect effects are not easily identified. In a nuclear power plant demonstration of STPA recently completed for the NRC, we found several ways that the non-safety system can impact the safety system indirectly and potentially lead to a catastrophe. In nuclear power design, major reliance is placed on the safety system to react when there is a disturbance in the plant and certification is usually focused on the safety system. Therefore, one of the most fundamental design requirements is that the successful operation of the safety system must be independent from the non-safety components of the plant. We demonstrated that the way the nuclear power community currently determines independence is flawed as it relies on identifying only direct relationships between the components and not indirect ones. The indirect relationships tend to be obscure and difficult to find. STPA can do this.

STECA has one other implication for SMS. Because STECA can be performed earlier than the traditional techniques, potential design flaws can be found earlier and eliminated from the design when more and cheaper design changes and options are available. This feature would appear to impact the current SMS design by allowing hazard analysis to begin earlier and to be more tightly intertwined with the original concept definition and architectural development rather than being only an after-the-fact analysis process. This advantage of STPA could save a lot of time and money in rework due to changes required late in the process.

12.   Conclusions

This report has presented a new approach, called systems-theoretic early concept analysis (STECA), for performing hazard analysis on a concept of operations and a safety-driven approach to concept development. STECA is based on systems- and control theory and its usefulness and practicality is demonstrated on an important aerospace application, called Trajectory-Based Operations.

Before describing STECA, Section 5 begins by describing general systems theory, upon which the rest of the report is built. Section 5 then presents a practitioner-oriented set of steps and heuristics for developing and analyzing models along with a rigorous theoretical development for academics and systems theorists seeking to advance the state of the art.

STECA is based on two basic steps. The first step involves recursively applying control-theoretic concepts using guide words, heuristics, and feedback control criteria to parse an existing ConOps document, resulting in the development of a hierarchical control model of the concept. The second step—analysis—consists of examining the resulting model with the explicit goals of identifying hazardous scenarios, information gaps, inconsistencies, and potential trade offs and

alternatives. That is, the analysis identifies incompleteness or gaps in the control structure, assures that all safety-related responsibilities are accounted for, and identifies sources of uncoordinated or inconsistent control.

The model development process itself provides a traceable database that can easily be queried and referenced back to the original concept documentation.

Section 6 demonstrates the principles developed in Section 5 on a real example concept being developed under the FAA's NextGen program. The model development and analysis yields important results, including several undocumented assumptions and scenarios associated with software behavior, potentially hazardous assumptions about human operator responsibilities, and potentially unsafe interactions among the various entities involved in TBO.

While Section 6 shows how the modeling effort and interrogation techniques can identify hazardous scenarios, Section 7 uses those results to document specific hazardous scenarios. These scenarios then drive the identification of safety-related requirements as well as the generation of control structure alternatives.

Comparing the TBO results from Sections 6 and 7 (and Appendix A) to existing analyses, performed by professional working groups using traditional techniques, Section 8 makes an assessment of STECA.

The primary research contributions are in integrating safety earlier in systems engineering activities (Figure 39) via rigorous, systematic methods.

System Engineering Phases
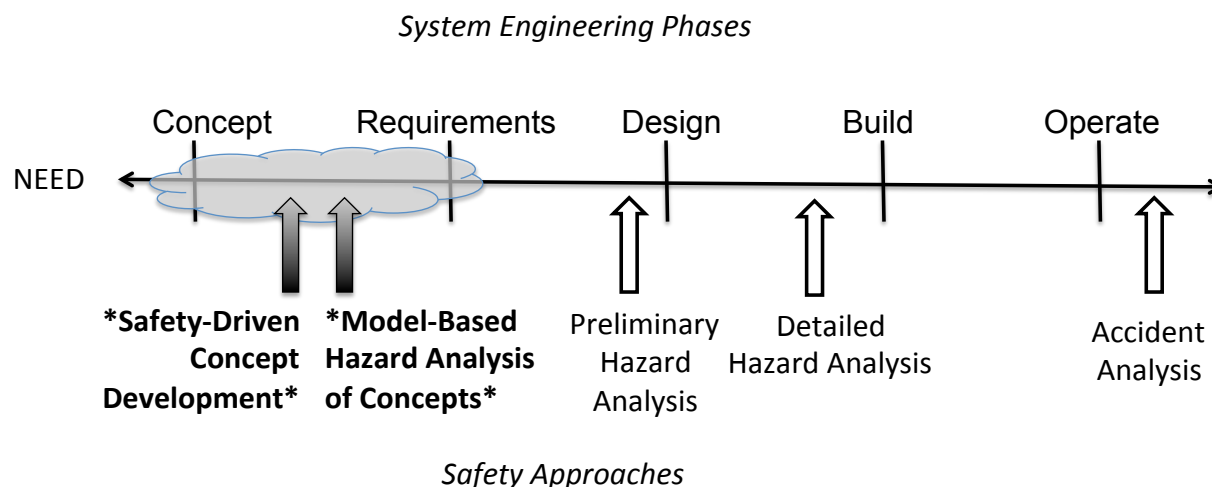
Figure 39. Research Contributions [adapted from Young, 2014]

The approach also fulfills many goals of model-based systems engineering (MBSE) and model based safety assessment (MBSA). That is, (1) the approach contains an unambiguous language based on control theory, (2) behavior is expressed in relationships that represent the "structure" of the system in terms of a hierarchical control model, and (3) the representation can be expressed from different views according to the system hierarchy. While the approach fulfills general goals of MBSE and MBSA, the safety-driven nature of the framework achieves specific objectives related to ensuring that safety is considered while the most important design and architecture decisions are made.

148

Specifically, STECA:

- Applies more rigor to the concept development process. The process described in Section 5 is repeatable and can be applied by individuals to understand a concept in systems-theoretic terms. Alternatively, the process can be used by teams and working groups to build consensus on how the system should (and should not) behave, to allocate responsibilities to various actors, and to define the interactions between those actors;
- Identifies a class of hazardous scenarios that existing techniques cannot find during concept development. Most accidents and incidents in modern, computer-intensive systems arise due to factors that extend beyond component reliability. Accidents arise due to unsafe interactions among components, which include software and human operator behavior, and STECA is more powerful than existing techniques in identifying these types of interactions;
- Makes explicit the assumptions that are often undocumented or implicit during early concept generation. Because ConOps documents are typically developed by subject matter experts, many details that are obvious to a particular expert may seem obvious and thus go undocumented. The model development approach in this report forces many of these assumptions to be made explicit, and often the various subject matter experts who generate the ConOps actually make competing or inconsistent assumptions about various aspects of the system. Section 6 demonstrates how some of these inconsistencies can be identified.

In addition to STECA, this report describes a more powerful approach to including human factors in hazard analysis. This new approach, still based on systems- and control theory, uses several important concepts from psychology and cognitive engineering in order to capture a larger, more rich set of causal factors that involve human behavior, automation, and the interaction of both. This new approach to human factors reveals causal factors that cannot be identified using traditional, failure-based hazard analysis techniques. Furthermore, the extended method reveals factors that were not found using the existing STAMP-based tools.

Finally, the report asserts that not only are STAMP, STPA, and STECA compatible with the existing framework, but also that the use of STAMP-based analytical tools satisfies certain objectives of SMS much more effectively that traditional tools. The use of STAMP, STPA, and/or STECA do not require a change to existing SMS processes. Moreover, these systems-theoretic approaches show the potential to improve and supplement existing activities such identifying mitigations and controls and for performing safety assessments earlier in the systems engineering process.

13.   Future Work

There are many potential paths of future research that build upon this work. While this report demonstrated how the results of the analyses can be used to generate alternative control structures, future work should demonstrate how these alternatives can be compared. Stakeholders identify potential tradeoffs or synergies, and tradeoffs could be made with respect to safety and/or extended to other system properties.

STECA was used in this report to generate and analyze a model of a ConOps. An interesting area of research relates to doing the inverse, or starting from model and generating a ConOps. The fact remains that many different stakeholders use a ConOps, not only engineers, and many of these stakeholders may not be comfortable with or fluent in the terminology and theory presented in the report. While the research presented here ultimately results in identifying scenarios and causal factors, those very scenarios can be used to refine the concept and, perhaps automatically, map into a more traditional format for Concept of Operations.

Tools should be generated to assist in the STECA process. Existing model-based systems engineering frameworks could be adapted or integrated into the systems- and control-theoretic processes described in Sections 5, 6, and 7.

Although this research focuses on the early phases of systems engineering (far left side of Figure 39), the framework has potential to applied to the very last phases of systems engineering (far right side of Figure 39). That is, when systems become operational, it is unfortunately necessary in some cases to perform accident and incident investigations. Accident reports typically share similar characteristics to Concept of Operations documents—they contain informal natural language text, use informal graphical depictions of events, and are often developed by committees comprised of potentially disparate views of the system. There exists an accident analysis process, called CAST (Casual Analysis using STAMP), based on the same systems- and control theory that has been successfully applied to accidents in a variety of domains [e.g. Dong, 2012; Hickey, 2012; Spencer, 2012; Hosse et al., 2013]. However, there is not yet a rigorous way to generate the necessary models from all the different sources of data associated with any major accident, and the methods presented in Section 5 represent a potential way to accomplish this goal.

With respect to human factors, while this report applied basic cognitive models to the human controller, more investigation into cognitive abstraction levels (e.g., Rasmussen's SRK Framework or Abstraction Hierarchy) should be pursued to account for the high-level goals and decisions that affect lower-level operational actions. If investigated, there might also be additional ways to categorize feedback and control actions in STAMP through an abstraction hierarchy as well. This may ultimately enhance the overall perspective on human control and lead to new ways of analyzing the human controller in STAMP with potential application to an automated controller.

14. Acronyms

Table 35. Acronyms

4DT	4-dimensional trajectory (3-dimensional positions and time-of-arrival at those positions)
ADS-B	Automatic Dependent Surveillance—Broadcast
ANSP	Air Navigation Service Provider
ATC	Air Traffic Control
ATM	Air Traffic Management

ConOps	Concept of Operations
DIM	Differential Importance Measure
ETA	Event Tree Analysis
FAA	Federal Aviation Administration (United States)
FMEA	Failure Modes and Effects Analysis
FTA	Fault Tree Analysis
GNSS	Global Navigation Satellite System
MAUT	Multi-Attribute Utility Theory
MBSE	Model Based System Engineering
MBSA	Model Based Safety Assessment
NAS	National Air Space
PHA	Preliminary Hazard Analysis
PHL	Preliminary Hazard List
RAW	Risk Achievement Worth
RNP	Required Navigation Performance
RTP	Required Time Performance
RRW	Risk Reduction Worth
STAMP	Systems-Theoretic Accident Model and Process
STECA	Systems-Theoretic Early Concept Analysis, based on STAMP accident model
STPA	Systems-Theoretic Process Analysis (hazard analysis method based on STAMP accident model)
TBO	Trajectory-Based Operations

15.   References

[Ale et al. 2008] Ale, B.J.M. and Bellamy, Linda and Cooke, Roger and Duyvis, Martina and Kurowicka, Dorota and Lin, Claire and Morales, Oswaldo and Roelen, Alfred, "Causal model for air transport safety", Modified Seventh Interim report (15 March 2007) (2008).

[Ashby 1957] Ashby, W Ross, An Introduction to Cybernetics (Chapman & Hall Ltd., 1957).

[ATSB 2013] ATSB, "Final Investigation Report AO-2010-089. In-flight uncontained engine failure Airbus A380-842, VH-OQA overhead Batam Island, Indonesia 4 November 2010, Australian Transportation Safety Board, 27 June", (2013).

[Ballin et al. 2008] Ballin, Mark G and Williams, David H and Allen, B Danette and Palmer, Michael T, "Prototype flight management capabilities to explore temporal RNP concepts", in Digital Avionics Systems Conference, 2008. DASC 2008. IEEE/AIAA 27th (2008), pp. 3--A.

[Barnett and Schulte 2001] Barnett, Mike and Schulte, Wolfram, "The ABCs of specification: AsmL, behavior, and components", (2001).

[Bass et al. 1998] Bass, Len and Clements, Paul and Kazman, Rick, Software Architecture in Practice, 2/E (Pearson Education India, 1998).

[Billings 1997] Billings, Charles E, Aviation automation: The search for a human-centered approach (1997).

[Boeing 2009] Boeing, "Statistical Summary of Commercial Jet Airplane Accidents", Commercial Airplanes, Worldwide Operations (2009).

[Boyd 1995] Boyd, J., *The Essence of Winning and Losing*, John Boyd Compendium, Project on Government Oversight: Defense and the National Interest, (1995). Available: http://pogoarchives.org/m/dni/john_boyd_compendium/essence_of_winning_losing.pdf

[Brown 2013] Brown, Laura, "Press Release - FAA Statement", (2013).

[Checkland 1999] Peter Checkland, Systems thinking, systems practice: includes a 30-year retrospective (John Wiley & Sons, Inc., 1999).

[Clements et al. 2003]Clements, Paul and Kazman, Rick and Klein, Mark, Evaluating software architectures (2003).

[Cone 2002]Cone, Edward, "The ugly history of tool development at the FAA", Baseline Magazine 4, 9 (2002).

[Cowlagi and Saleh 2013] Cowlagi, Raghvendra V and Saleh, Joseph H, "Coordinability and consistency in accident causation and prevention: formal system theoretic concepts for safety in multilevel systems", Risk analysis 33, 3 (2013), pp. 420-433.

[Crawley et al. 2004] Edward Crawley and Olivier de Weck and Steven Eppinger and Christopher Magee and Joel Moses and Warren Seering and Joel Schindall and David Wallace and Daniel Whitney, "The Influence of Architecture in Engineering Systems", in 1st Engineering Systems Symposium (2004).

[Crawley et al. 2008] Crawley, Frank and Preston, Malcolm and Tyler, Brian, HAZOP: Guide to best practice: guidelines to best practice for the process and chemical industries (IChemE, 2008).

[Croft 2005] Croft, John, "Airbus and Boeing spar for middleweight title", Aerospace America 43, 7 (2005), pp. 36-38.

[Daouk and Leveson June 2001] Mirna Daouk and Nancy G. Leveson, "An Approach to Human-Centered Design", in Workshop on Human Error and System Development, Linkoping, Sweden (June 2001).

[de Weck 2009] de Weck, Olivier, "Fundamentals of Systems Engineering", (Massachusetts Institute of Technology: MIT OpenCouseWare), http://ocw.mit.edu (Accessed 7 August 2013). License: Creative Commons BY-NC-SA (2009).

[de Weck et al. 2011] de Weck, Olivier L and Roos, Daniel and Magee, Christopher L, Engineering systems: Meeting human needs in a complex technological world (The MIT Press, 2011).

[Dekker 2005] Dekker, Sidney, Ten questions about human error: A new view of human factors and system safety (CRC Press, 2005).

[Dekker 2013] Dekker, Sidney, The Field Guide to Understanding Human Error: Second Edtion (Ashgate Publishing, Ltd., 2013).

[DoD 2012] US DoD, "MIL-STD-882E, Department of Defense Standard Practice System Safety" (2012).

[Dong 2012] Dong, Airong, "Application of CAST and STPA to railroad safety in China" (2012).

[Drake et al. 2010] Drake, Bret G and Hoffman, Stephen J and Beaty, David W, "Human Exploration of Mars, Design Reference Architecture 5.0", in Aerospace Conference, 2010 IEEE (2010), pp. 1-24.

[Dulac and Leveson 2005] Dulac, Nicolas and Leveson, Nancy, "Incorporating safety in early system architecture trade studies", in International System Safety Conference Proceedings (2005).

[Dumoulin 1988] Dumoulin, Jim, NSTS 1988 News Reference Manual (Kennedy Space Center, NASA., 1988).

[Dunjó et al. 2010] Dunjó, Jordi and Fthenakis, Vasilis and V́ılchez, Juan A and Arnaldos, Josep, "Hazard and operability (HAZOP) analysis. A literature review", Journal of hazardous materials 173, 1 (2010), pp. 19-32.

[Eppinger and Browning 2012] Eppinger, Steven D and Browning, Tyson R, Design structure matrix methods and applications (MIT Press, 2012).

[Ericson 2005] Clifton A. Ericson, Hazard Analysis Techniques for System Safety (John Wiley & Sons, Inc., Hoboken, NJ, USA, 2005).

[FAA 1998] "Order: Safety Risk Management," FAA, Order 8040.4, Jun. 1998.

[FAA 2007] "Advisory Circular: Introduction to Safety Management Systems (SMS) for Airport Operators," FAA, AC 150/5200-37, Feb. 2007.

[FAA 2008] FAA, "Safety Management System Manual" (2008).

[FAA 2010a] "Safety Management System (SMS): Recommendations and Final Report," FAA Aviation Rulemaking Committee, Mar. 2010.

[FAA 2010b] "SAFETY MANAGEMENT SYSTEM (SMS) FRAMEWORK," FAA, Jun. 2010.

[FAA 2011a] "Design And Manufacturing (D&M) Safety Management System (SMS) Pilot Project Guide," FAA, Rev C, Oct. 2011.

[FAA 2011b] "Developmental Guidance for Design And Manufacturing Safety Management System Framework," FAA, Nov. 2011.

[FAA 2012a] "Order: Aviation Safety (AVS) Safety Management System Requirements," FAA, VS 8000.367A, Nov. 2012.

[FAA 2012b] J. Duven, "2012 Manufacturers Safety Management System Pilot Project Report Design and Manufacturing Organizations," FAA, 2012.

[FAA 2012c] FAA, "NextGen Segment Implementation Plan 5.0", Federal Aviation Administration, Washington, DC (2012).

[FAA 2013a] "Order: Safety Management System," FAA, VS 8000.369A, May 2013.

[FAA 2013b] "Order: Aviation Safety (AVS) Safety Policy," FAA, VS 8000.370A, Apr. 2013.

[FAA 2013c] FAA, "NextGen Implementation Plan", Washington, DC, Available: http://www.faa.gov/nextgen/implementation/media/NextGen_Implementation_Plan_2013.pdf (accessed March 18, 2014) (2013).

[FAA 2014a] "The Safety Management System (SMS) Manual version 4.0," FAA, 2014.

[FAA 2014b] M. Reinert and W. Desrosier, "Report from Part 21/Safety Management Systems (SMS) Aviation Rulemaking Committee to the FAA: Recommendations on Certification Procedures for Products and Parts," FAA, Washington D.C., Oct. 2014.

[FAA 2015] "Advisory Circular: AC 120-92B Safety Management Systems for Aviation Service Providers," FAA, Washington DC, AC 120-92B, Jan. 2015.

[Fed. Regist. 2010] "Notice of Proposed Rulemaking: Safety Management System for Certificated Airports," Fed. Regist., vol. 75, no. 194, p. 62008, Oct. 2010.

[Fed. Regist. 2015] "Safety Management Systems for Domestic, Flag, and Supplemental Operations Certificate Holders," Fed. Regist., vol. 80, no. 5, pp. 1308–1328, Jan. 2015.

[Figueira et al. 2005] Figueira, José and Greco, Salvatore and Ehrgott, Matthias, Multiple criteria decision analysis: state of the art surveys vol. 78, (Springer, 2005).

[Fingerhut et al. 2010] Fingerhut, Benjamin P. and Zinth, Wolfgang and de Vivie-Riedle, Regina, "The detailed balance limit of photochemical energy conversion", Phys. Chem. Chem. Phys. 12 (2010), pp. 422-432.

[Flach 2004] Flach, John, A cognitive approach to situation awareness: theory and application (Ashgate Pub Limited, 2004).

[Flach and Voorhorst 2012] Flach, J.M. and F. Voorhorst, *What Matters*, Unpublished manuscript, Dayton, OH, (2012). Available: http://psych-scholar.wright.edu/flach/publications

[Fleming et al, 2012] C. H. Fleming, M. Spencer, N. Leveson, and C. Wilkinson, "Safety Assurance in NextGen," NASA, Langley, VA, NASA/CR–2012-217553, Mar. 2012.

[Fleming et al. 2012] Fleming, Cody and Ishimatsu, Takuto and Miyamoto, Yuko and Nakao, Haruka and Katahira, Masa and Hoshino, Nobuyuki and Thomas, John and Leveson, Nancy, "Safety Guided Spacecraft Design using Model-Based Specifications", in Proceedings of the 5th IAASS Conference (2012).

[Fleming et al. 2013] Fleming, Cody Harrison and Spencer, Melissa and Thomas, John and Leveson, Nancy and Wilkinson, Chris, "Safety assurance in NextGen and complex transportation systems", Safety Science, Elsevier 55 (2013), pp. 173-187.

[Forrester 1987] Forrester, Jay W, "Nonlinearity in high-order models of social systems", European Journal of Operational Research 30, 2 (1987), pp. 104-109.

[Friedenthal et al. 2007] Friedenthal, Sanford and Griego, Regina and Sampson, Mark, "INCOSE model based systems engineering (MBSE) initiative", in INCOSE 2007 Symposium (2007).

[Frola and Miller 1984] Frola, FR and Miller, CO, "System safety in aircraft management", Logistics Management Institute, Washington DC (1984).

[Harel 1987] Harel, David, "Statecharts: A visual formalism for complex systems", Science of computer programming 8, 3 (1987), pp. 231-274.

[Harkleroad et al. 2013] Eric Harkleroad and Adan Vela and James Kuchar and Bob Barnett and Rose Merchant-Bennett, "ATC-045 Risk-based Modeling to Support NextGen Concept Assessment and Validation", MIT Lincoln Laboratory & Federal Aviation Administration (2013).

[Hickey 2012] Hickey, Jon Jon Patrick, "A system theoretic safety analysis of US Coast Guard aviation mishap involving CG-6505" (2012).

[Hollnagel et al. 2007] Hollnagel, Erik and Woods, David D and Leveson, Nancy, Resilience engineering: Concepts and precepts (Ashgate Publishing, Ltd., 2007).

[Hosse et al. 2013] Hosse, RenÃ and Spiegel, Dirk and Welte, Jan and Schnieder, Eckehard and others, "Integration of Petri Nets into STAMP/CAST on the Example of Wenzhou 7.23 Accident", in Control and Automation Theory for Transportation Applications vol. 1, no. 1 (, 2013), pp. 65-70.

[ICAO 2013a] "Annex 19 to the Convention on International Civil Aviation: Safety Management," ICAO, Jul. 2013.

[ICAO 2013b] "Safety Management Manual (SMM) 3rd Edition," ICAO, Doc 9859, 2013.

[ICAO 2013c] "Global Aviation Safety Plan," ICAO, Montreal, Quebec, Doc 10004, 2013.

[ICAOd] "ANNEX 6 – Operation of Aircraft (Amdt 33-B)," ICAO.

[INCOSE 2011] INCOSE, SE, "INCOSE Systems Engineering Handbook v. 3.2. 2", INCOSE-TP-2003-002-03.2. 2. October (2011).

[Irvine 2007] Dean Irvine, "Shooting for the moon: The new space race, (accessed) (August) 2013", (2007).

[Ishimatsu et al. 2010] Ishimatsu, Takuto and Leveson, Nancy and Thomas, John and Katahira, Masa and Miyamoto, Yuko and Nakao, Haruka, "Modeling and hazard analysis using STPA", in Conference of the International Association for the Advancement of Space Safety, Huntsville, Alabama (2010).

[Jackson 2010] Jackson, Michael RC, "Role of avionics in trajectory-based operations", Aerospace and Electronic Systems Magazine, IEEE 25, 7 (2010), pp. 12-19.

[Jackson et al. 2009] Jackson, MRC and Gonda, J and Mead, R and Saccone, G, "The 4D trajectory data link (4DTRAD) service-Closing the loop for air traffic control", in Integrated Communications, Navigation and Surveillance Conference, 2009. ICNS'09. (, 2009), pp. 1-10.

[Jones 1990] Jones, Cliff B, "Systematic software development using VDM", Prentice-Hall (1990).

[JPDO 2010] JPDO, "Joint Planning and Development Office Concept of Operations for the Next Generation Air Transportation System", Version 3.2 (September 30, 2010) (2010).

[JPDO 2011] JPDO, "JPDO Trajectory-Based Operations (TBO) Study Team Report", Joint Planning and Development Office (2011).

[JPDO 2012] JPDO, "Capability Safety Assessment of Trajectory Based Operations v1.1", Joint Planning and Development Office Capability Safety Assessment Team (2012).

[Kapurch 2010] Kapurch, Stephen J, NASA Systems Engineering Handbook (DIANE Publishing, 2010).

[Keeney et al. 1979] Keeney, Ralph L and Raiffa, Howard and Rajala, David W, "Decisions with multiple objectives: Preferences and value trade-offs", Systems, Man and Cybernetics, IEEE Transactions on 9, 7 (1979), pp. 403-403.

[Klien et al. 2004] Klien, G and Woods, David D and Bradshaw, Jeffrey M and Hoffman, Robert R and Feltovich, Paul J, "Ten challenges for making automation a a ``Team Player" in Joint Human-Agent Activity", Intelligent Systems, IEEE 19, 6 (2004), pp. 91-95.

[Kollewe and Gabbatt 2010] Julia Kollewe and Adam Gabbatt, "Qantas to replace half of its Rolls-Royce A380 superjumbo engines (accessed August 2013)", (2010).

[Leonard et al. 2007] Leonard, John and Barrett, David and How, Jonathan and Teller, Seth and Antone, Matt and Campbell, Stefan and Epstein, Alex and Fiore, Gaston and Fletcher, Luke and Frazzoli, Emilio and others, "Team MIT urban challenge technical report", (2007).

[Leveson 1995] Nancy G. Leveson, Safeware: System Safety and Computers (Addison Wesley, 1995).

[Leveson 2000a] Leveson, Nancy, "Completeness in formal specification language design for process-control systems", in Proceedings of the third workshop on Formal methods in software practice (2000), pp. 75-87.

[Leveson 2000b] Leveson, Nancy G., "Intent specifications: An approach to building human-centered specifications", Software Engineering, IEEE Transactions on 26, 1 (2000), pp. 15--35.

[Leveson 2004] Leveson, Nancy, "A new accident model for engineering safer systems", Safety Science 42, 4 (2004), pp. 237-270.

[Leveson 2009] Leveson, Nancy G, "Software challenges in achieving space safety", (2009).

[Leveson 2012] Nancy G. Leveson, Engineering a Safer World (MIT Press, 2012).

[Leveson and Reese 1994] Leveson, NG and Reese, JD, "TCAS II intent specification" (1994).

[Leveson et al. December 2001] Nancy Leveson and Maxime de Villepin and Mirna Daouk and John Bellingham and Jayakanth Srinivasan and Natasha Neogi and Ed Bachelder and Nadine Pilon and Geraldine Flynn, "A Safety and Human-Centered Approach to Developing New Air Traffic Management Tools", in ATM 2001, Albuquerque NM (December 2001).

[Lutz and Carmen Mikulski 2003] Lutz, Robyn R and Carmen Mikulski, Ines, "Operational anomalies as a cause of safety-critical requirements evolution", Journal of Systems and Software 65, 2 (2003), pp. 155-161.

[McCurdy 1993] McCurdy, Howard E, "Inside NASA: High technology and organizational change in the US space program", New Series in NASA History, Baltimore: Johns Hopkins University Press,| c1993 1 (1993).

[Mesarovic 1970] Mesarovic, Mihajlo D, "Multilevel systems and concepts in process control", Proceedings of the IEEE 58, 1 (1970), pp. 111-125.

[Mesarovic and Takahara 1975] Mesarovic, Mihajlo D and Takahara, Yasuhiko, General systems theory: mathematical foundations vol. 113, (Academic Press New York, 1975).

[Mesarovic et al. 1970] Mesarovic, MD and Macko, D and Takahara, Y, "Theory of Multilevel Hierarchical Systems", New York, Academic (1970).

[Morari and Stephanopoulos 1980] Morari, Manfred and Stephanopoulos, George, "Studies in the synreport of Control Structures for Chemical Processes: Part II: Structural aspects and the synreport of alternative feasible control schemes", AIChE Journal 26, 2 (1980), pp. 232-246.

[Morari et al. 1980] Morari, Manfred and Arkun, Yaman and Stephanopoulos, George, "Studies in the synreport of control structures for chemical processes: Part I: Formulation of the problem. Process decomposition and the classification of the control tasks. Analysis of the optimizing control structures", AIChE Journal 26, 2 (1980), pp. 220-232.

[Nakao et al. 2012] Nakao, H and Katahira, M and Miyamoto, Y and Leveson, N, "Safety Guided Design of Crew Return Vehicle in Concept Design Phase Using STAMP/STPA", in ESA Special Publication vol. 699, (2012).

[Netjasov and Janic 2008] Netjasov, Fedja and Janic, Milan, "A review of research on risk and safety modelling in civil aviation", Journal of Air Transport Management 14, 4 (2008), pp. 213-220.

[Office 1986] United States. General Accounting Office, Air traffic control: FAA's advanced automation system acquisition strategy is risky: report to the Secretary of Transportation (The Office, 1986).

[Owre et al. 1996] Owre, Sam and Rajan, Sreeranga and Rushby, John M and Shankar, Natarajan and Srivas, Mandayam, "PVS: Combining specification, proof checking, and model checking", in Computer Aided Verification (, 1996), pp. 411-414.

[Patteau 2009] Patteau, Eric, "SESAR Joint Undertaking - Annual Report 2009", (2009).

[Pereira et al. 2006] Pereira, Steven J and Lee, Grady and Howard, Jeffrey, "A system-theoretic hazard analysis methodology for a non-advocate safety assessment of the ballistic missile defense system", DTIC Document (2006).

[Petty 2002] Petty, John Ira, "Inflight Crew Escape System", (2002).

[Rasmussen 1975] Rasmussen, Norman C, Reactor safety study: An assessment of accident risks in US commercial nuclear power plants vol. 4, (1975).

[Rasmussen 1983] Rasmussen, J., *Skills, Rules, and Knowledge; Signals, Signs, and Symbols, and Other Distinctions in Human Performance Models*, IEEE Transactions on Systems, Man and Cybernetics, Vol. SMC-13, No. 3, (1983).

[Rasmussen 1986] Rasmussen, J., "Information Processing and Human Machine Interaction: An Approach to Cognitive Engineering. 1986", New York. North-Holland 1, 2 (1986), pp. 3.

[Rasmussen 1997] Rasmussen, Jens, "Risk management in a dynamic society: a modelling problem", Safety science 27, 2 (1997), pp. 183-213.

[Rasmussen and Goodstein 1987] Rasmussen, Jens and Goodstein, LP, "Decision support in supervisory control of high-risk industrial systems", Automatica 23, 5 (1987), pp. 663-671.

[Ray 2014] Ray, Elizabeth, "FAA Order 7110.65 V Air Traffic Control. (April, 2014)", Federal Aviation Administration, Washington, DC: Air Traffic Rules and Procedures Service (2014).

[Reason 1990] Reason, James, Human error (Cambridge university press, 1990).

[Reason 2000] Reason, James, "Human error: models and management", BMJ: British Medical Journal 320, 7237 (2000), pp. 768.

[Reason 2008] Reason, James T, The human contribution: unsafe acts, accidents and heroic recoveries (Ashgate Publishing, 2008).

[Roland and Moriarty 2009] Harold E. Roland and Brian Moriarty, Preliminary Hazard Analysis, in System Safety Engineering and Management, Second Edition (John Wiley & Sons, Inc., Hoboken, NJ, USA, 2009).

[RTCA 2008] RTCA, "Safety, Performance and Interoperability Requirements Document for the In-Trail Procedure in the Oceanic Airspace (ATSA-ITP) Application", DO-312, Washington DC (2008).

[RTCA 2011] RTCA, "Safety, Performance and Interoperability Requirements Document for Airborne Spacing Flight Deck Interval Management (ASPA-FIM)", DO-328, June 19 (2011).

[SAE 1996] SAE, "ARP-4761, Guidelines and methods for conducting the safety assessment process on civil airborne systems and equipment" (1996).

[Simon 1977] Simon, Herbert A, "The structure of ill-structured problems", in Models of discovery (Springer, 1977), pp. 304-325.

[Skogestad 2004] Skogestad, Sigurd, "Control structure design for complete chemical plants", Computers & Chemical Engineering 28, 1 (2004), pp. 219-234.

[Spencer 2012] Spencer, Melissa Beth, "Engineering financial safety: a system-theoretic case study from the financial crisis" (2012).

[Sterman 1994] Sterman, John D, "Learning in and about complex systems", System Dynamics Review 10, 2-3 (1994), pp. 291-330.

[Stieglitz 1948] Stieglitz, William I, "Engineering for safety", Aeronautical Engineering Review 7, 2 (1948), pp. 18-23.

[Strafaci 2008] Strafaci, Adam, "What does BIM mean for civil engineers?", CE News, Tranportation (2008).

[Takahara and Mesarović 2003] Takahara, Yasuhiko and Mesarović, Mihajlo D, Organization structure: cybernetic systems foundation vol. 22, (Springer, 2003).

[Tatjewski 2008] Tatjewski, Piotr, "Advanced control and on-line process optimization in multilayer structures", Annual Reviews in Control 32, 1 (2008), pp. 71-85.

[Thomas 2013] Thomas, John, "Extending and automating a systems-theoretic hazard analysis for requirements generation and analysis", Extending and Automating a Systems-Theoretic Hazard Analysis for Requirements Generation and Analysis. SAND2012-4080, Sandia National Labroratories (2013).

[Thornberry 2014] Thornberry, Cameron. Extending the Human-Controller Methodology in Systems- Theoretic Process Analysis (STPA), MIT Master's Thesis June (2014).

[Vesely et al. 1981] Vesely, William E and Goldberg, Francine F and Roberts, Norman H and Haasl, David F, "Fault tree handbook", DTIC Document (1981).

[Vicente 1999] Vicente, Kim J, Cognitive work analysis: Toward safe, productive, and healthy computer-based work (CRC Press, 1999).

[Vincoli 2005] Jeffrey W. Vincoli, Basic Guide to System Safety, Second Edition (John Wiley & Sons, Inc., Hoboken, NJ, USA, 2005).

[Woods et al. 2010] Woods, David D and Johannesen, Leila J and Cook, Richard I and Sarter, Nadine B, Behind human error (Ashgate Publishing Company, 2010).

[Yeo 2010Yeo, Ghim-Lay, "A380 flights to Los Angeles unprofitable with Trent 900s: Qantas (accessed August 2013)", (2010).

[Yeo 2013] Yeo, Ghim-Lay, "FAA approves 787 battery changes (accessed August 2013)", (2013).

[Young 2014] William Young, "STPA-SEC for Cyber Security / Mission Assurance" (2014).

[Young and Leveson 2014] Young, William and Leveson, Nancy G, "An integrated approach to safety and security based on systems theory", Communications of the ACM 57, 2 (2014), pp. 31-35.

[Zheng et al. 1999] Zheng, Alex and Mahajanam, Rama V and Douglas, JM, "Hierarchical procedure for plantwide control system synreport", AIChE Journal 45, 6 (1999), pp. 1255-1265.

Appendices

A. Case Study - TBO Conformance Monitoring

This appendix continues the analysis from Section 6. That section focused on the ground-based component, while the following analysis focuses on the airborne component. Following the airborne conformance monitoring analysis, this appendix concludes with a formal presentation of "Analyzing Safety-Related Responsibilities" and "Coordination and Consistency" that was started in Section 6.

A.1   Airborne Conformance Monitoring Model

The process model for the ground-based control agent (i.e. the ANSP) is simply,

$$\mathbf{PM}_G := \{\mathbf{Conf}_i, \dot{x}_o\} \tag{49}$$

where $\mathbf{Conf}_o$ represents whether the aircraft, o for "ownship", is conforming with its prescribed trajectory, and $\dot{x}_o$ represents some prediction of the future states of the ownship aircraft. As with the analysis of ground-based systems, the flight deck-based control actions are unspecified in the TBO ConOps' section regarding conformance monitoring. It is also assumed here that the airborne control algorithm is a function of conformance and that the piloting function will generate an action that enforces conformance.

The conformance model proceeds in a similar fashion as above (equations 47–50). The airborne conformance monitor is a mapping

$$\mathcal{B}_{\mathcal{L}A} := \mathcal{V}_{\mathbf{mA}} \rightarrow \mathcal{I}_{\mathbf{Ac}} \tag{50}$$

$\mathcal{I}_{Ac}$ is the signal going to the piloting function (i.e. the airborne control agent) regarding conformance, and the functional behavior of the sensor is

$$\mathcal{V}_{mA} = \mathcal{L}_A \times D_{c,o} \tag{51}$$

where $\mathcal{L}_A$ is a model of the airspace state and $D_{c,o}$ is the decision criteria regarding conformance of the ownship aircraft o. The "airborne", (or piloting function) conformance model is defined as the set of dynamic variables,

$$\mathcal{L}_A := \{\mathbf{z}_{\mathbf{int},o}, \mathbf{z}_{\mathbf{act},o}, \rho, T, P_r, W, E_{cm}, F_D\} \tag{52}$$

where

$$z_{\mathbf{int},0} := \{G, C, t\}_{\mathbf{int},o}$$
$$z_{\mathbf{act},0} := \{G, C, t\}_{\mathbf{act},o}$$
$$\rho = \mathbf{Traffic\ Density}$$
$$\tau = \mathbf{Operation\ Type}$$
$$P_r = \{\mathbf{RNP}, \mathbf{RTP}\}$$
$$W = \mathbf{Wake\ turbulence\ model}$$

$$E_{cm} \coloneqq \textbf{Elliptical conformance model}$$
$$F_D \coloneqq \{F, z_{\textbf{int},o}\}$$

with similar definitions as equation (49) but with the ownship subscript o. Criteria for determining whether the ownship aircraft conforms with its assigned trajectory are,

$$D_{co} = \{z_{\textbf{act},0} | z_{\textbf{act},o} \notin \overline{Z_o}(z_{\textbf{int},o}, E_{cm}, a_A)\} \tag{53}$$

where $\overline{Z_o}$ is an allowed volume for ownship aircraft o, as a function of the intended aircraft state in time, the elliptical conformance model at a given time (E_{cm}), and an alert parameter set by the airborne operator. Like equation 50, evaluation of equation (56) to True indicates that the ownship aircraft o does not conform to its assigned trajectory.

The control algorithm of the airborne control agent (flight crew) consists of—at a minimum— use of some decision on conformance alerting. That is,

$$\textbf{CA}_A \supseteq \textbf{Conf}_o \tag{54}$$

where $\textbf{Conf}_o$ is obtained, according to equation 9, via signal processing of the input signal $\mathcal{I}_{Ac}$.

A.2 System-Level Conformance Monitoring Model

The following model (Figure 40) represents all of the development in the previous section, as well as the main body (Section 6.3). This model is a continuation of the conformance monitoring model developed in the main body but includes relevant details from the rest of the TBO ConOps. Relevant details to Conformance Monitoring include the types of actions necessary to close or maintain a trajectory.
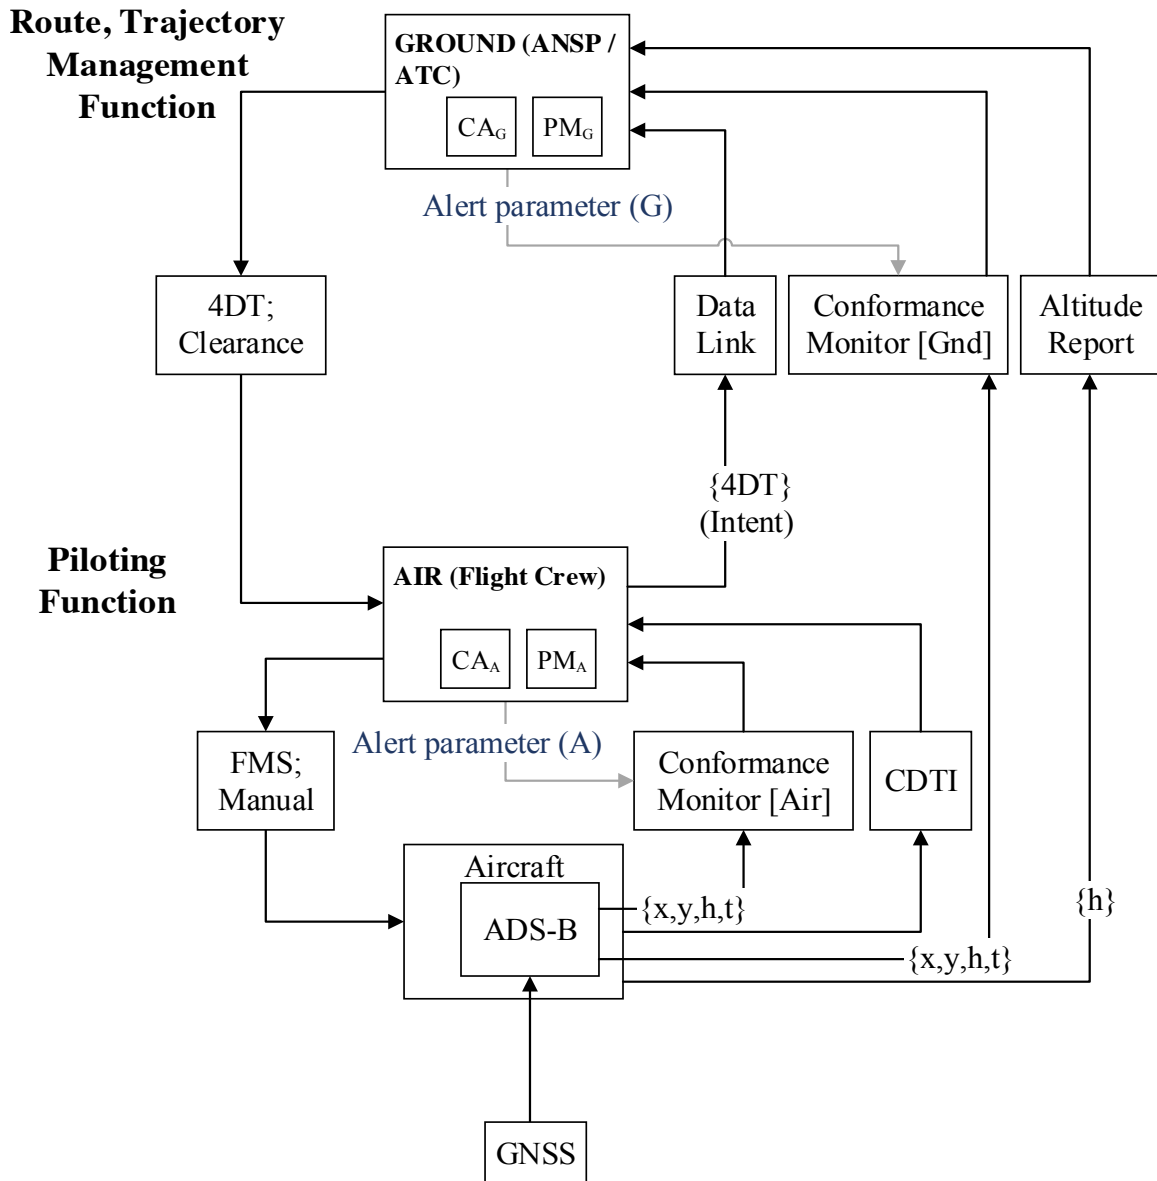
Figure 40. General Conformance Monitoring System Model

A.3   *Airborne Conformance Monitoring Analysis*

This section is a continuation of the model development and analysis in Sections 6.3 and 6.3.1 of the main body.
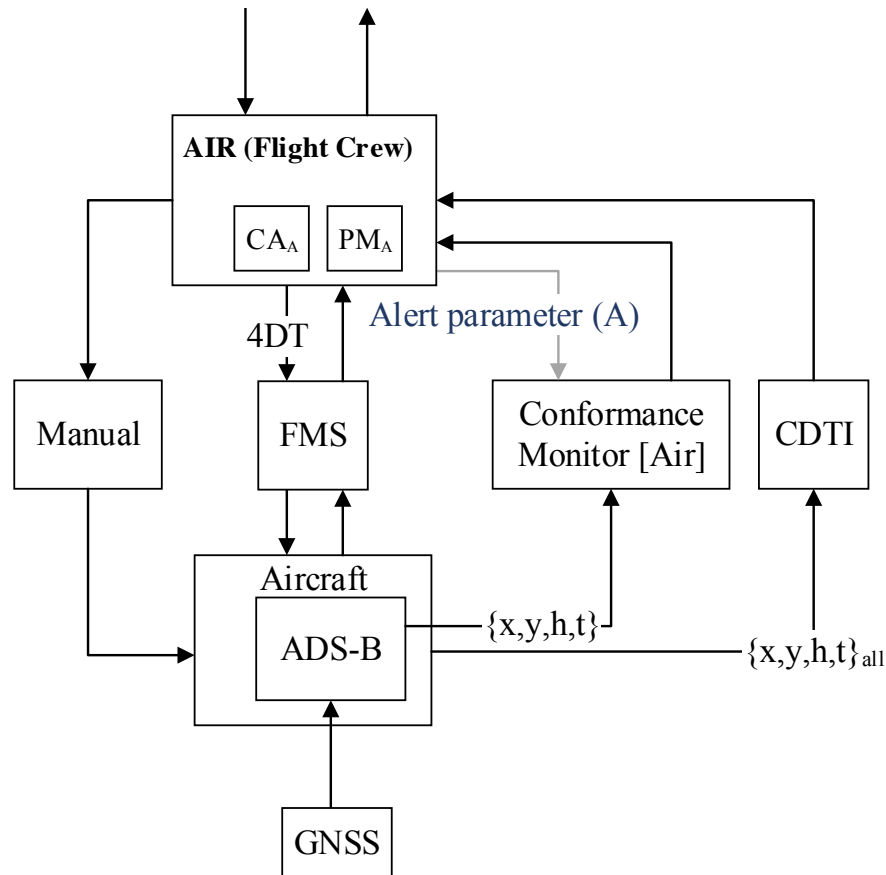
Figure 41. Flight Deck (Airborne) Control Loops

- *Goal Condition* - what are the goal conditions, and how can they violate safety constraints and safety responsibilities?

Similar to the ANSP analysis above (elsewhere), the Flight Crew is also responsible for merging, sequencing, and spacing, as well as assuring conformance. Most of the discussion in the ANSP is applicable here, because the Flight Crew and ANSP must work together to achieve flow and spacing goals and to implement achievable trajectories.

In terms of conformance monitoring and the goal of assuring conformance, the TBO ConOps states that the "pilot must also work to close the trajectory" if the aircraft does not conform [JPDO, 2011, p.14]. The goal of conforming—closing a trajectory—should never take precedence over crew or passenger safety. Of course there are scenarios where conformance will be necessary for ensuring safety, but any assumption that conformance will always ensure safety is tenuous. For example, if there is severe turbulence and the TBO automation cannot update the trajectory, then it should certainly not be a goal of the pilots to close the trajectory. The high level

164

requirements should thus state that a pilot should not work to achieve[32] the desired trajectory if it causes loss of separation, controlled flight into terrain, flight into restricted areas, loss of aircraft control, or any of the system-level hazards.

- *Action Condition* - how does the controller affect the state of the system? Are the actuators adequate or appropriate given the process dynamics?

  The flight crew affects the state of the system by loading trajectory parameters into the Flight Management System (FMS). Elsewhere, i.e. outside of the section on conformance monitoring, the TBO ConOps acknowledges that a 4DT could be flown manually, with limited performance. The technical specification of a 4DT is relatively ill-defined, and thus it is unclear how this action occurs other than to "load" the information into the FMS [e.g. JPDO, 2011, pp.11,42,58]. Other concepts provide higher fidelity information about FMS and the parameters of a 4DT [e.g. Ballin et al., 2008; Jackson et al., 2009], although issues still remain with respect to completeness of this control loop.

  The ConOps prescribes the use of the FMS to "close" the trajectory. In terms of the Action Condition, there are several issues with this approach. Not least is the fact that the FMS is the primary resource used to maintain the 4D trajectory; it follows that one cause of non-conformance is due to performance of the FMS itself. In such a scenario, the very thing that is causing the goal to be missed is then required to correct it. Furthermore, the crew will be required to input new parameters to enter into the FMS in order to close the trajectory. It is unclear how the flight crew will calculate these parameters. Again, the ConOps relies on TBO automation for many of these calculations, but it is (possibly) this very automation that caused non-conformance to begin with. Finally, it is expected that flight crews will sometimes be required to fly manually, but the TBO ConOps itself acknowledges that flight crews cannot be expected to fly with sufficient accuracy in order to achieve some of the 4DT performance parameters.

  For the merging, sequencing, and spacing Goal Condition (and given current regulatory standards), the action condition is only marginally satisfied. Currently, an aircraft can only control its own trajectory via FMS or manual inputs and does not have the authority to issue instructions to other aircraft. Merging, sequencing, and spacing depend on the actions of every aircraft entering the airspace or traffic flow. In the current system, spacing is ensured in part by collision avoidance systems, but these systems are part of a larger policy that implements actions on all aircraft involved in a conflict. Likewise, merging and sequencing must be part of a larger policy that coordinates actions among multiple aircraft. Such a policy is not

---

[32] This discussion involves *achieving* or closing a trajectory, however similar requirements would be developed about *maintaining* or following a trajectory.

necessarily in conflict with the TBO concept, but it is not explicitly defined. The TBO ConOps mentions "tools" that will help pilots monitor other aircraft (see below, Observability Condition), but it does not define action conditions that should result from the use of these tools.

- *Model Condition* - what states of the process must the controller ascertain? How are those states related or coupled dynamically? How does the process evolve?

  With respect to conformance, the Flight Crew analysis has similar results to the ANSP analysis. In particular, the Flight Crew will have similar human factors issues to the ANSP if they simply maintain a binary model of conformance. See section 6.4 for analysis of the ANSP Model Condition.

  For merging, sequencing, and spacing, the crew (and associated automation) must maintain not only a model of ownship state but also all the other aircraft entering its vicinity. This model should contain information about the current state of all the aircraft along with their intent. Because a flight crew and on-board automation do not have authority to issue instructions to other aircraft, the model must also be updated to include any changes in pilot intent, new clearances from ANSP, and non-conformance of other aircraft in the space. It is these latter factors that are relatively undefined in the ConOps, and thus the Model Condition for the airborne control loop is not completely satisfied.

  Given that the Flight Crew will still have ultimate responsibility over the aircraft[33], there are further human factors issues with assigning Merging, Sequencing, and Spacing responsibility. Problems with the Action Condition have already been described, but it is also unclear how Flight Crews will maintain a model of both the aircraft itself (and all its complexity) along with the state and intent of all the other aircraft. The Observability Condition analysis below describes one proposed approach to this problem.

- Observability *Condition* - how does the controller ascertain the state of the system? Are the sensors adequate or appropriate given the process dynamics?

  The observability analysis for the Flight Crew is similar to that of the ANSP elsewhere. From the perspective of the flight crew and its goal of maintaining closed trajectories, conformance monitoring alone does not satisfy the Observability Condition. Like the ANSP, the flight crew must have information about the degree to which the aircraft is not conforming, which requires more than an alerting function. Many of the early proposals to implement concepts related to TBO involve

---

[33] The TBO ConOps advocates for continuing current regulation with respect to crew responsibility: "the flight crew's authority over the aircraft's trajectory (FAR 91.3) do not change with trajectory negotiations" [JPDO, 2011].

time-based metering. In time-based metering, conformance is relative to a scheduled time of arrival, and feedback simply involves comparing the estimated time of arrival to scheduled time of arrival and then adjusting the speed accordingly. As 4DT operations involve more dimensions, including increasingly tight lateral and vertical constraints, the feedback will become more complex.

For merging, sequencing, and spacing, the airborne component must have access to all aircraft that will be part of the traffic flow or entering into the ownship trajectory. As proposed, this will be done via Cockpit Display of Traffic Information (CDTI). While CDTI systems will display current state information for all the aircraft in a given vicinity, what is not clear in terms of the flight crew's Goal Condition is how the crew will discern the intent of all the aircraft. That is, all the predicted trajectories may not be directly observable via a CDTI, nor will recent, current, or future clearances or trajectory changes. All of these conditions affect the ability of a flight crew to achieve merging, sequencing, and spacing objectives.

The analysis of the ANSP control conditions concluded with a brief discussion about role allocation, interfaces, and other architectural considerations. Assuming that pilots have some separation authority, and they almost certainly will still have authority over their own aircraft per regulation, similar factors to the ANSP analysis apply here.

A.4 Analyzing the Safety-Related Responsibilities

The analysis in Section 6.4 describes the systems-thinking that goes into identifying potential flaws and scenarios associated with the ConOps. The following analysis describes the underlying mathematical formalism that can be used to more rigorously identify flaws and missing information.

Hazard - Aircraft violate minimum separation

Generally, loss of separation occurs whenever the protected airspace of any two aircraft overlap. The loss of separation hazard is thus:

$$\mathcal{H}_s = \left\{ \mathfrak{C}, \breve{z} \,\middle|\, \mathfrak{C}_i \times z_i \mapsto \breve{z}_i, \mathfrak{C}_j \times z_j \mapsto \breve{z}_j, \breve{z}_i \cap \breve{z}_j \neq \emptyset \right\} \tag{55}$$

where $\breve{z}_k$ is a protected airspace volume of aircraft k around the current state $z_k = (x_k(t), y_k(t), h_k(t))$ and x, y, h represent latitude, longitude, and altitude, respectively. Traditionally this has been done via a "hockey puck" model, where $\breve{z}$ is a cylinder with 5NM diameter and height of 1000' with the aircraft at its center. Loss of separation occurs when any two of these virtual cylinders intersects. The TBO ConOps proposes other models of protected airspace, which will be discussed shortly.

Assume now that the responsibility of conformance is fulfilled; that is, an action is applied to maintain or resume conformance. This conformance condition can be defined as:

$$\mathcal{M}_c = \left\{ \mathfrak{C}_k, V_k \,\middle|\, \mathfrak{C}_k \times z_k \mapsto z'_k, z'_k \in V_k \right\} \tag{56}$$

167

where V_k is a conformance volume of aircraft k around the desired state $\check{z}_k = (\check{x}_k(t), \check{y}_k(t), \check{h}_k(t))$. Finally, conformance monitoring sounds an alert whenever an aircraft deviates from the desired state. See the development of equations (47) and (55) elsewhere, and observe that the aircraft state, z_k, is a function of ground track and climb performance in that development. Does conformance alerting enforce safety constraints? The answer is yes and no.

Reiterating the definition from Section 5, let $\mathbf{P}(x, \mathbb{S})$ be defined for all pairs $(x, \mathbb{S})$, where

$$\mathbf{P}(x, \mathbb{S}) \equiv x \text{ yields } \mathbb{S}. \tag{57}$$

The predicate $\mathbf{P}(x, \mathbb{S})$ is true, iff $\mathbb{S}$ is a defined condition or system state, and x is an action resulting in that condition. A safety-related conflict occurs if the following proposition does not hold:

$$\neg \exists c \in \mathfrak{C}[\mathbf{P}(c, \mathcal{H}_s) \wedge \mathbf{P}(c, \mathcal{M}_c)] \tag{58}$$

That is, there is a conflict with safety responsibilities if there is an action that can simultaneously result in the loss of separation hazard, $\mathcal{H}_s$, and fulfill the conformance condition, $\mathcal{M}_c$. Such an action is possible if there are any aircraft (or any other debris or hazardous situation) in the presence of the intended aircraft space, V_k.

Any argument against the proposition in (61) must also make the relatively strong assertion that there will never be any conflict along the protected trajectory of an aircraft. An obvious rebuttal to this argument might go as follows. If aircraft α is not conforming and must "work to close" trajectory, it is equally plausible that aircraft β is in the same situation. It also plausible that aircraft β is now on the very trajectory that α is attempting to regain. The goal of assuring conformance must be tempered by the larger goal of assuring separation.

Hazard - Aircraft enters uncontrolled state

Although there could be any number of reasons an aircraft loses control, consider the following formal definition of the hazard.

$$\mathcal{H}_c = \{\mathfrak{C}, \bar{z}_i | \mathfrak{C}_i \times z_i \mapsto z_i, z_i \cap \bar{z}_i \neq \emptyset\} \tag{59}$$

where $\bar{z}_i$ is the allowable flight envelope of aircraft i. That is, a loss of control hazard occurs when the aircraft exceeds its flight envelope. This definition omits common causes of loss of control, but the omission is intentional. Because this analysis focuses on trajectory-based operations, the definition in equation (62) includes scenarios where the trajectory *itself* causes loss of control. There are myriad scenarios leading to loss of control, and analysis of these scenarios constitutes a lower level in the system hierarchy. Given the definition of conformance in equation (59), there is a potential conflict if the following does not hold:

$$\neg \exists c \in \mathfrak{C}[\mathbf{P}(c, \mathcal{H}_c) \wedge \mathbf{P}(c, \mathcal{M}_c)] \tag{60}$$

There is a potential conflict with safety responsibilities if an action simultaneously exceeds aircraft capability, causing loss of control $\mathcal{H}_c$, and fulfills the conformance condition, $\mathcal{M}_c$. Such

an action is possible if the trajectory generation algorithm is inadequate or if it has inaccurate information about aircraft capability, aircraft status, and environmental conditions. The TBO ConOps thoroughly documents the necessary parameters for computing aircraft capability and including this in TBO automation.

A.5 Coordination and Consistency

Conformance is intended to assure consistency among the various actors in TBO, including the ANSP, flight crews, and operating centers. Despite its intent, it is actually a source of potential inconsistencies and lack of coordination.

Consider again a conformance model (generalized from the development elsewhere and included again here for reference). Conformance monitoring and alerting is a mapping

$$\mathcal{B}_{cm} \coloneqq \mathcal{L}_{cm} \times D_{cm} \to \mathcal{I}_{cm},$$ (61)

where $\mathcal{L}_{cm}$ is a model of the airspace state and D_{cm} is the decision criteria regarding conformance.

$$\mathcal{L}_{cm} \coloneqq \{z_{int}, z_{act}, \rho, T, P_r, W, E_{cm}, F_D\}$$ (62)

$$z_{int} \coloneqq \{G, C, t\}_{int}$$
$$z_{act} \coloneqq \{G, C, t\}_{act}$$
$$\rho = \textbf{Traffic Density}$$
$$\tau = \textbf{Operation Type}$$
$$P_r = \{\text{RNP}, \text{RTP}\}$$
$$W = \textbf{Wake turbulence model}$$
$$E_{cm} = \textbf{Elliptical conformance model}$$
$$F_D = \{F, z_{int}\}$$

$$D_{cm} = \{z_{act} | z_{act} \notin \breve{z}(z_{int}, E_{cm}, a_{cm})\}$$ (63)

where $\breve{z}$ is an allowed volume for an aircraft. Conformance alerting is a function of current surveillance data in four dimensions and desired aircraft state in four dimensions, some allowed tolerance volume, and an "Alert Parameter" ($z_{act}, z_{int}, E_{cm}, a_{cm}$, respectively, in equation 66).

This mathematical formulation of conformance monitoring and alerting helps establish issues with coordination and consistency, in at least three ways. First, the mapping, $\mathcal{L}_{cm} \times D_{cm}$ is a function of an "Alert Parameter", a_{cm}. This parameter is available to any agent with a conformance monitor, ground or airborne, and is thus independent and potentially inconsistent. For example, the ground controller sets an alert parameter for his or her own monitor, while the flight crew independently does so for their monitor. The TBO ConOps does not describe or specify the rationale for including this function, but it may be assumed that it is to counteract alarm overload or over- and under-sensitivity. Furthermore, the TBO ConOps does not specify what the alert parameter entails with respect to design. The ConOps does refer to existing aircraft

functions such as altitude alerts for the airborne monitor, but it is unclear how either the ground or airborne agents will "set" these parameters.

A.6 Analysis

In addition to the coordination scenarios described in Section 6.4, there is another coordination problem from the perspective of the ANSP (Figure 21). This issue arises because the ANSP receives conformance information about every aircraft, $i \in N$. Given Aircraft A and B in Figure 22, assume that Aircraft B is non-conforming. Aircraft B is early, which represents a potential conflict because it will arrive at a crossing route with Aircraft A. While their intended trajectories do not conflict, it has become a coordination problem, particularly if Aircraft B cannot conform quickly. That is, the ANSP has to coordinate the actions given to A and B (and potentially other up- and downstream aircraft) in order to avoid a conflict. The TBO ConOps recognizes such a problem and suggests the use of tactical clearances by air traffic controllers along with collision avoidance systems.

B.   TBO Model Development and Analysis

This appendix demonstrates how the methods described in Section 5 can be used to store the model in a database and ensure traceability to the original, natural language text. The model is stored in a different format than the format used in Section 6 and Appendix A Appendices

Case Study - TBO Conformance Monitoring. This appendix also represents a different aspect of TBO—negotiation of 4D trajectories. The graphical model (Figure 25) is included in the main body in Section 7.

The following tables have six columns: (1) text quoted directly from the TBO ConOps [JPDO, 2011]; (2)–(5) are each element of the tuple $(\mathcal{S}, \mathcal{R}, \mathcal{B}, \mathcal{A})$, respectively; and (6) is a description of the updated control model. Abbreviations are used for various elements of the control model, e.g. "CA" means "Control Algorithm".

Original Text	Source	Role	Behavior	Context	Control Model
Quote directly from TBO ConOps	$\mathcal{S}$	$\mathcal{R} \in \{\mathcal{C}, \mathcal{K} \ldots \mathcal{P}, \mathcal{L}\}$	$\mathcal{B}_c, \mathcal{B}_{\mathcal{K}} \ldots \mathcal{B}_{\downarrow}p, \mathcal{B}_{\downarrow}\mathcal{L}\}$	$\mathcal{A}$:=Context, assumptions, or rational for choice of model elements	Model Elements Abbreviations CA:=Control Algorithm PM:=Process Model

Table 36. Analysis for Negotiation with ANSP

Original Text	Source	Role	Behavior	Context	Control Model
From the ANSP perspective, the main emphasis will be on planning and enhanced system prediction,					
with the objective of making most strategic ATM actions prior to flight departure or, where this is not	ANSP	Controller	Generates actions		CA (policy for strategic versus tactical ['tactical' used later])
practical, before the flight is forecast to enter a block of airspace or an airport terminal area where a					PM (block of airspace, terminal, constraints); CA (determining when to act relatve to knowledge in PM)
constraint exists. The ANSP moves towards ``Management by Exception,'' maximizing collaborative					CA ('collaborative')
pre-flight iterative planning, and refining the proposed trajectory using all available relevant					CA (what is considered 'pre-flight' versus what is not)
information as the takeoff time approaches, minimizing the need for in-flight intervention. The SESAR					PM (take off time); CA (policy for what is considered 'take off time approaches')

It is likely that most problems, whether flow management involving many aircraft or separation	Aircraft & Airspace	Controlled Process	Interacts	Separation and flow	PM
management involving two aircraft, will be solved as early as practical, as soon as the information					Process (2 or more conflicting aircraft, flow of a/c)
about the problem is good enough to support deriving an acceptable solution, with low probability that					PM and CA (PM: 'goodness' of data; CA to determine what to do with relative 'goodnesses')
the ``problem'' was a false alert, or that the ``solution'' fails to solve the problem. This is to give the best					
chance of minimizing deviations from the operator's desired trajectory and to promote overall					CA (objectives: stability, predictability, early negotiation)
trajectory stability and predictability. Thus, the ANSP will initiate negotiation earlier rather than later.	ANSP	Controller	Generates actions		

The ANSP will generally try to make minor trajectory changes through timing or speed adjustments	ANSP	Controller	Generates actions		Cntl Algorithm (Policy for preferring minor traj changes); Action (Speed adjustments)
where possible for separation, initiating airborne merging and spacing or other aircraft procedures, and					CA (merging / sequencing algorithm); PM (separation, spacing requirements, merging)
flow management. Path modification would be used when timing adjustments are insufficient or not					CA (decision about insufficiency of spd/tm adjustment); Action (path modification)
the best solution. The overriding objective will be to maintain trajectory stability (and thus the ANSP					Cntl Algorithm (cntl objectives, policy for maintaining traj stability)
prediction functions) and minimize the need for tactical vectoring, but how this is achieved will depend	ANSP	Controller	Generates actions		CA (objectives: Policy for minimizing tactical maneuvers)
on the conditions and type of operation. Minor trajectory updates will replace most of today's open-					CA (decision on 'minor' traj mods versus open vector)

174

loop vectors for spacing or separation. Vectors are not really a negotiation in that the flight crew will	Flight Crew	Controller	PM		Control Action (Vector); Safety resp (fc must follow clearance); Control Action (Execute heading change, [FC])
accept the clearance unless they have an overriding safety reason to not do so. Vectors take the aircraft					PM (Open v closed traj)
onto an open trajectory that subsequently must be closed with a new 4DT.					
P.14-15					
The pilot must also work to close the trajectory. Pilots will need to update waypoints leading to a	Pilot	Controller	Actions		Control Action (update FMS/waypoints)
closed trajectory in the FMS, and work to follow the timing constraints by flying speed controls.					PM (waypoints/times of a closed Trajectory); Actuator (FMS); Cntl Action (speed controls)

175

Under dense traffic conditions, buffers will be needed in the system so that local trajectory changes can					PM of ANSP (traffic density); CA of ANSP (4DT or other for ensure buffer ->policy a function of density, desired stability/propagation)
be made without propagating beyond the aircraft initially affected. ANSP automation will predict	ANSP Auto	Sensor		lower level cntlr, possible a "sensor" - algorithm to predict aircraft traj given conds and capabilities.	PM (constraints, ac type, weather, ac traj)
conditions when it will become infeasible for aircraft to meet future constraints, and will propose	Human ATC	Controller	form PM	see next line	PM (Aircraft capability, predicted future trajectory [via automation], alt trajs [ibid])
trajectory updates to the controller, possibly before the aircraft start requesting them, both to optimize	Human ATC	Controller	form PM		CA (policy of preempting aircraft traj requests)
overall system flows and to reduce communications. For instance, if winds are more strongly out of the					CA (objectives: Policy for reducing communication)
west than forecast at Dallas-Ft. Worth (DFW), all the flights from the west coast will be arriving early,					PM (weather, trajectory prediction)

and those from the east coast will be arriving late. This will trigger revised times at the arrival fixes.					
Users may include predefined preferences (similar to the FMS cost index function) with the filed flight	Users	Sensor		(actually external info)	Input / Feeback (?) (Cost index; predefined preferences; flight plan)
plan for ANSP automation to consider in proposing alternative trajectories.					Controller Ouput (not necessarily a cntl action: alt trajectories)
The ANSP may negotiate by proposing revised constraints to the aircraft or the FOC, and let the	ANSP	Controller	Generates actions; transmits info		PM (necessity of revised constraints [further questions]); CA (policy for determining whether to send rev constr, alt trajs, or new [single] traj)
operator choose a preferred trajectory to meet the constraints. This new preferred trajectory could then	Operator	Controller	Generates actions	This is a lower level controller than ANSP, but above FC	PM (preferred traj. Does 'pref traj' = '4DT' [i.e. 'actual traj'])
become the authorized trajectory that now has to be renegotiated if a further change is required. This is					

177

likely to be the case when the ANSP is using the trajectory for separation management.	Trajectories	Actuator			Action (Traj for separation mgmt)
Performance requirements for airspace and operations will generally be set before flight, but	Performance requirement	Actuator	Translation		PM (performance reqs; time of application of perf reqs)
occasionally might be used opportunistically during flight. This might occur, for instance, during					
en route convective weather, where the ANSP negotiates optimal routes for high-performing aircraft					PM (convective weather, ac performance)
(data link, low RNP) through gaps in the weather, and lower performing aircraft are routed around it.					PM (RNP, communication type, weather/gaps); CA (algorithm for determining optimal vs route around)
In this example, the better-equipped aircraft gets the advantage of using the more direct route.					CA (policy for routes / equippage)

Table 37. Analysis for Negotiation with FOC

Original Text	Source	Role	Behavior	Context	Control Model
The dispatcher is typically a person responsible for managing flights; one or more dispatchers perform	Dispatcher	Controller	Generates actions		Controller (Dispatcher)
their work in a FOC. For single-aircraft operations, the flight planning function may be performed by	FOC	Controlled Process			Controller (Collection of Dispatchers); Action (Flight Planning); Actuator (Flight Plan)
the pilot in command, member of the flight crew, or another entity on their behalf such as a flight					Controller (PIC, FC, Flight Service Station, 3rd Party Co); PM (responsible flight planner)
service station, or to a third-party company who offers collaborative air traffic management (CATM)					
service to Part 91 operators.					
The flight planner's horizon for negotiating a flight may vary from tens of minutes to weeks. Typical	FOC	Controller			PM (time horizon); Alg (when to make decision)

reasons for negotiating a trajectory with the ANSP include the following:					``Reasons'' –- control algorithm and process model
- For an airline, build and update a schedule for its overall flight operation to meet the schedule					Process (Schedule); Action (Build/update overall flight operation)
- Get best initial trajectory for individual aircraft before takeoff					CA (Trajectory ['best' –- CI]); Process (Individual Aircraft)
- Prioritize among multiple flights (under the flight planner's jurisdiction) entering congested					CA ([FOC] Prioritization policy); PM (multiple flights; airspace congestion)
airspace or terminal area					
- Re-route, delay or substitute one or multiple flights around weather or congestion to maintain	Trajectory	Actuator			CA ([FOC] Prioritization policy, re-route/delay/substitute policy); PM (multiple flights; airspace congestion; weather)
business objectives					

- Diversion from original planned destination due to severe conditions, weather, fuel needs,					CA ([FOC] Prioritization policy, re-route/delay/substitute policy; diversion); PM (multiple flights; airspace congestion; weather); Proc Disturbance (severe condition, weather, fuel need, aircraft emergency, passenger consideration)
aircraft emergency, passenger considerations, etc.					
The FOC will maintain a model of the aircraft performance and other key characteristics to anticipate	FOC				PM (a/c performance, key characteristics, expected flight progress, environment, fuel usage)
the expected flight progress in the given environment, fuel usage, etc. This information informs the					
flight planner of what trajectory options are feasible for the aircraft when negotiating a new trajectory					PM (feasible trajectories); Alg (negotiation); Cntl Input (from ANSP)
with the ANSP. New processes and protocols by which revised trajectories are negotiated and					PM/Act/Sens/Alg (processes and protocols);

approved between the FOC, ANSP, and aircraft may be needed. This may include some form of partly					Cntl Input (ANSP); Process (a/c)
automated negotiation to replace the daily and hourly teleconferences between the ANSP and various	Partly automated negotiation (ANSP, FOC)	Controller (or sensor)			
FOCs to strategize about weather and other events.					PM (weather, other events); Alg ('strategize')
Preflight planning and flight following are key roles of the FOC in order to develop and maintain the	FOC	Controlled Process	``flight following'' is the controlled process; preflight planning is a sensor		Process (FOC following flights)
business plan and business trajectory of the operator through optimization of both the individual					Alg (optimization); PM (individual a/c; fleet); Action (Business Plan, Business Trajectory)
aircraft and the fleet. This includes specification of the airframe to be used to conduct the operation,					Alg (spec of ac, fuel, FC); PM (capability of a/c, FC; fuel(?))

fuel decisions, and flight crew assignments. Once payload and fuel decisions have been made and the					Alg (spec of ac, fuel, FC, payload); PM (capability of a/c, FC; fuel, payload)
fuel for the flight has been loaded, flexibility is very limited. This is especially true for very long haul					PM (time of fuel loading; route [long v short-haul]; aircraft weight); Alg (haul length & fuel, loading)
flights limited by weight. Typically, these decisions are made anywhere from a few hours before the					CA (determining time [rel to departure] of decision)
flight up to the time of departure, depending on the latest payload, weather, and other related					PM (time of dept, payload, weather forecast)
information. As with the ANSP, the objective is to make most strategic decisions before the flight					Alg (decision time [pre- v post-departure])
commences. Even during the flight, the dispatcher is the primary and preferred decision maker for					PM (negotiator [FOC, FC]); Alg (decision on who to negotiate with)

strategic negotiations with the ANSP because the FOC has access to more information, and the					
negotiation can take place over net-centric operations. The cockpit is also part of net-centric operations	Cockpit (Crew or Automation)	Controlled Process AND Controller			
and works with their dispatcher in concurrence on changes. The FOC will generally be negotiating	Net-centric ops	Sensor, actuator			
trajectories greater than 20 to 30 minutes into the future with the ANSP, and their role in negotiation	ANSP	Controller			PM (time to "negotiation even"); Alg (decision on who to negotiate with; 20-30 minute criterion)
diminishes relative to the flight deck as the time gets closer to the event for which the negotiation was					
initiated. Once a revised trajectory is negotiated, this new trajectory is conveyed to the pilot for	Pilot	Controlled Process	Translate higher level action into lower level one		PM (revised trajectory [for all controllers – FOC, FC, ANSP])

approval and execution. The expanded role for the FOC enabled by the shift to more strategic decision-					CA (execution); Cntl Status (approval)
The ANSP provides a forum to facilitate collaboration between flight planners representing multiple	Flight planner	Controller		Flight planner is then in 'control' of the multiple FOCs	
FOCs when there is a system-wide event or constraint. Among multiple operators, aggregate solutions	System-wide event	Controlled Process		Event is actually a disturbance to CP	Process Disturbance (system-wide event or constraint)
to demand/capacity imbalances may be proposed to and by the ANSP. They should provide improved					PM (Demand, Capacity); CA (comparing imbalance --); Action (``aggregate solutions'')
operations within the context of the operator's business objectives in comparison to solutions that					PM (Operator business objective); CA (policy function of biz obj, demand, capacity)
might be individually imposed by the ANSP. Common situational awareness across the flight planning	Common situational awareness	Sensor	Continuous, discrete data	Assume net-centric operations based on previous data	

participants improves the options for dealing with constraints.					

Table 38. Analysis for Negotiation with Pilots

Original Text [JPDO, 2011]	Source	Role	Behavior	Context	Control Model
During the flight, the flight crew complies with the cleared trajectory except in emergencies; non-	Flight crew	Controlled Process	Translate	In Emergencies it becomes 'top level' controller	Process (FC - compliance w/ traj)
emergency changes are negotiated and agreed to before being executed. Flight crews will have access					PM (emergency status [y/n or continuum])
to 4D weather and NAS status information relevant to their flight through network-centric operations.	Network-centric Ops	Sensor	Provide binary or cont info	Sensor is actually something else, but net-centric provides the ``lines'' between blocks	PM (F/C PM: 4D weather; NAS status; both ``relevant to flight'')

186

Some aircraft may have sophisticated flight-planning functionality onboard, including trial planning	Flight Planning Avionics	Sensor	Provide binary or discrete info	This is an intelligent sensor; could provide binary 'Y/N' for trajectories as well as discrete parts of flight plan(s)	Sensor (Flt Pln Avionics - trail plan, eval of trajs); PM (F/C PM: 'eval of trajs')
and evaluation of proposed trajectories. However, even with advanced airborne decision-support	Advanced Airborne Decision-support Automation	Sensor	Provide binary or discrete info	See above comment	
automation, pilot-initiated trajectory negotiations may be limited by workload considerations.	Flight crew	Controlled Process		Part of the FC output is also in terms of feedback, in this case negotiation or emergency	PM (grnd: neg from F/C)
The pilot monitors progress toward meeting assigned constraints and initiates negotiation directly or	Flight crew	Controller	Form, update process model		PM (Assigned constraints, progress to constraints; negotiating entity (FC/FOC))
through the FOC if projected to be unable to meet a constraint. If the aircraft can still meet the					PM (Able/unable to meet constraints)

constraints of the 4DT, but would prefer to renegotiate the constraint for efficiency or scheduling					CA (efficiency or scheduling); Actuator (Constraint)
reasons, the pilot or FOC may request negotiation of the constraint, or the preferred constraint changes					
may be listed as alternatives in the flight plan. When a pilot requests a minor trajectory change that still	Flight crew	Controller	Generate action	Request is only a precurser to actually generating action on aircraft	CA (policy for determining type of change); PM (delta-Traj meets/doesn't meet constraints; status of request); FB to higher cntl (change request)
meets all constraints, this trajectory change should be a fairly easy process expedited by local ANSP	Local ANSP Automation			Unclear if this automation has decision authority; thus unclear what it's role is. ``Expedited'' implies that the automation may simply approve it. This statement also implies that constraints are not the only thing that defines a trajectory	Actuator (Trajectory; Constraints)

automation. Negotiating a change in constraints can be much more complex, since constraints have	Constraint negotiation	Actuator			
typically been negotiated and agreed upon through collaborative air traffic management (CATM)	CATM Procedure	Controller	Generate action	This is in the general ANSP controller function	CA (negotiation, agreement criteria)
procedures.					
For most airspace and operations, trajectories will include some flexibility (sometimes referred to as	Trajectories	Actuator		Trajectories are the constraints that ultimately define aircraft behavior, if followed	Actuator (Trajectory; Constraints; Windows); PM ([both grnd and FC]: traj, consts; windows)
``windows") that allows the operator to optimize within limits without renegotiating the trajectory.	Operator	Controller	Generate action		CA (optimization alg); PM (cost index/function; limits of window)
This might apply to an aircraft choosing a path in real time between thunderstorms. Conformance					

monitoring on the ground considers the range of these windows, so minor maneuvering is authorized	Conformance Monitoring	Sensor	Provide binary and cont info		Sensor (Aircraft state; Conformance notification); PM ([grnd] Window; Conformance; 4D Contract); Actuator (4D Contract)
within the context of the 4D contract.					
In 2025, aircraft will vary widely in their ability to accurately adhere to a 4DT and in their ability to					
exchange trajectory information with the ground. At the lower end of performance, some aircraft are	Aircraft	Controlled Process			Process (Capability, Information Exchange)
only capable of adhering to wide lateral trajectories and flexible timing requirements (windows), while	Aircraft	Controlled Process			Actuator (Trajectory; Constraints; Windows); PM ([both grnd and FC]: traj, consts; windows; aircraft capability); Process (aircraft)
high-performing aircraft can transmit via data link detailed 4D trajectories that can be executed with	High-performing aircraft			Datalink is the line between boxes of control model, not an actual sensor/actuator	Sensor (Datalink air-to-grnd)

high accuracy. In general, NAS operations and ANSP decision-support automation must be designed	NAS Operations				
to deal with the entire spectrum of aircraft capabilities. However, significant operational efficiencies	ANSP decision-support automation				PM (Aircraft capability); CA (Efficiencies; policy w/r/t aircraft capability)
can be gained by taking advantage of the additional information-sharing capabilities and performance	High-performing aircraft	Controlled Process	Interact with env		Process (Navigation performance; Info-sharing capabilities)
accuracies of high-performing aircraft.					
The far-term JPDO Operational Improvements include a self-separation capability, which may be	Self-separation	Controller		This implies the pilot/ac as its own control agent	
restricted to airspace where only self-separating aircraft can operate or may include mixed-equipage					
environments where some aircraft are self-separating while others are ANSP-managed . There are also					

a number of pair-wise delegated separation operations in mid- and far-term NextGen, such as airborne	Delegated-separation	Controller			
merging and spacing, parallel runway operations, and oceanic procedures. Self-separating aircraft	Merging and spacing, parallel runway ops, oceanic procedures	Controlled Process		Controlled Process (Operation type)	
separate themselves using ADS-B as the primary means of airborne surveillance. ADS-B provides the	ADS-B	Sensor	Transmit continuous and discrete data	Discrete data is aircraft plan/trajectory; continuous is aircraft state	Sensor (ADS-B [for self-sep aircraft cntl loop])
location and velocity vector, plus the near-term intent of other self-separating aircraft in the vicinity.	ADS-B	Sensor	Transmit continuous and discrete data	Discrete data is aircraft plan/trajectory; continuous is aircraft state	PM (Location, velocity, near-term intent, self-separating aircraft, 'vicinity')
Research is required to determine how to safely and effectively utilize the capabilities of these					
advanced aircraft in increasing efficiency and throughput of NextGen operations. When an aircraft has					

been delegated some separation responsibility; there needs to be sufficient flexibility in the trajectory					Actuator (Trajectory with flexibility)
that maneuvers for separation can be accomplished without trajectory negotiation. In this case, the					Action ([FC] Manuever inside of [ANSP-sanctioned] trajectory)
trajectory might be constrained only sufficiently to support traffic flow management and would be					PM ([ANSP] constraints); CA (traffic flow mgmt - constraints sufficiency/flexibility tradeoff)
virtually useless for separation management decisions. However, these self-separating aircraft are					CA (decision wrt usefulness of constraints for flow vs separation mgmt)
themselves operating based on highly accurate trajectories, which are continuously maintained to be					PM (accuracy of trajectory)
conflict free using onboard strategic and tactical conflict detection and resolution functionality.	Onboard avionics			Role allocation is subject to discussion (CD&R can be modeled as a controller or a sensor for pilots)	

To ensure predictable operations, trajectory changes that result in the near-term generation of new	Trajectory change	Actuator			CA (control objective: predictability, no near-term conflicts), PM ('near-term'), Action (trajectory change)
conflicts are not acceptable. To sustain situational awareness, the new trajectory path is broadcast via					PM (new trajectory path)
ADS-B before the aircraft begins maneuvering. This is especially important outside of the range of any					
ANSP surveillance. If this highly accurate trajectory were made available to the ground via data link,	ANSP Surveillance	Sensor	Continuous data	This implies that ANSP does NOT use ADS-B as a primary data source	Sensor (ANSP surveillance)
then the ANSP decision-support automation could accurately predict the movement of both self-	ANSP decision-support automation				PM (prediction of movement; self-, delegated- and ANSP-separated aircraft)
separating and delegated-separation aircraft. If future ANSP separation management decision-support					

194

automation were designed to effectively use this more accurate information for conflict detection and				
resolution functions, then more effective use of airspace capacity should result. Thus, an open research				
question is whether there should, in fact, be two trajectories in effect for these aircraft: a negotiated	Trajectories	Actuator			PM ([ANSP] Negotiated vs ``4D Intent'' Trajectory); CA ([ANSP] deciding on when/if to issue different version of traj)
trajectory used for traffic flow management that provides sufficient flexibility for self-separation				
maneuvers, and a highly accurate ``4DT intent'' used for separation management that is periodically				
updated by the aircraft without negotiation. NAS operational efficiency would probably benefit from				

having access to 4DT intent information from all aircraft that are capable of generating, transmitting					Action (performing 'highly accurate' 4DT); CA (generating 'highly accurate' 4DT); PM ([ANSP] All a/c – 'highly accurate' 4DT)
via data link, and performing to a highly accurate 4DT.					Sensor ([a/c - ANSP] datalink and 'highly accurate' 4DT)

C. Formalism of Hazardous Control Actions

The following definition of a hazardous (unsafe) control action comes from the work of Thomas [2013]. A hazardous control action in the STAMP accident model can be expressed formally as a 4-tuple $(C^i, T, \mathbb{C}, Co)$ where:

C^i is the source controller that can issue control actions in the system. The controller may be automated or human.

T is the type of control action. There are two possible types: Provided describes a control action that is issued by the controller while Not Provided describes a control action that is not issued.

$\mathbb{C}$ is the control action (i.e. command) that is output by the controller.

Co is the context in which the control action is or is not provided.

Each element of an unsafe control action is a member of a larger set. The following properties must hold:

(1) $C^i \in S$, where S is the set of source controllers in the system (C^i is then the source of action in each level of the system, per equations (19) and (22))
(2) $T \in \mathcal{T}$, where $\mathcal{T}$ = {Provided, Not Provided}
(3) $\mathbb{C} \in \mathcal{O}_a(C^i)$ where $\mathcal{O}_a(C^i)$ is the set of control actions that can be provided by controller C^i
(4) $Co \in \rho(C^i)$, where $\rho(C^i)$ is the set of potential contexts, or process model states, for controller C^i

To assist in enumerating or aggregating individual contexts, the context Co is further decomposed into variables, values, and conditions:

(1) V is a variable or attribute in the system or environment that may take on two or more values. For example, heading, position, and airspeed are three potential variables for a aircraft.
(2) VL is a value that can be assumed by a variable. For example, M0.85 is a value that can be assumed by the variable airspeed.
(3) Cd is a condition expressed as a single variable/value pair. For example, aircraft conformance is a condition.
(4) The context Co is the combination of one or more conditions and defines a unique state of the system or environment in which a control action may be given.

The following additional properties related to the context of a hazardous control action can therefore be defined:

(5) $V \in \mathcal{V}(C^i)$, where $\mathcal{V}(C^i)$ is the set of variables relevant to the system hazards $\mathcal{H}$
(6) $VL \in \mathcal{VL}(V)$,, where $\mathcal{VL}(V)$ is the set of values that can be assumed by variable V
(7) $Cd = (V, VL) \in \mathcal{CD}(C^i)$, where $\mathcal{CD}(C^i)$ is the possible set of conditions for controller C^i
(8) $Co = (Co_1, Co_2, \ldots)$, where each C o i is independent. That is, no two Co_i refer to the same variable V.

Finally, each hazardous control action must be linked to a system-level hazard:

(9) To qualify as a hazardous control action, the action $(\mathcal{C}^i, T, \mathbb{C}, Co)$ must be able to cause a hazard $H \in H$, where H is the set of system level hazards.

A hazardous control action expressed as a 4-tuple $(\mathcal{C}^i, T, \mathbb{C}, Co)$ must satisfy the above properties 1-9.